# Contents

alten 400





2

# Contents

# Contents

Printer





3

# Contents

# Contents

# Contents

# Contents





5

# Contents

6

# Contents

# Figures



# Tables





7

# Module IN — Introduction

## Objectives

Upon completion of this module, you will be able to:

- Identify responsibilities of a system administrator.
- Identify the HP-UX documentation targeted towards helping system administrators.
- Discuss the contents of this course.

## IN-1. SLIDE: Welcome

**Welcome**

**Welcome to the**

**HP-UX System Administration Course**

51436 IN-1. 1 

### Student Notes

There is a lot of knowledge required to effectively administer an HP-UX system. This course is designed to teach you enough information to become an effective administrator, and to teach you how to learn more on your own.

This student guide is designed to be used both in class and as a reference tool. There is a Table of Contents at the front of the student guide that outlines all the topics. Each topic page has a **SLIDE** and additional information entitled **Student Notes**. If you have time, read through the Student Notes before each lecture. This will help you get the most out of your classroom time. After lecture, use the Student Notes to reinforce and review what you have learned.

## IN-1. SLIDE: Welcome

Welcome

Welcome to the

HP-UX System Administration Course

51436 IN-1. 1 

### Student Notes

There is a lot of knowledge required to effectively administer an HP-UX system. This course is designed to teach you enough information to become an effective administrator, and to teach you how to learn more on your own.

This student guide is designed to be used both in class and as a reference tool. There is a Table of Contents at the front of the student guide that outlines all the topics. Each topic page has a **SLIDE** and additional information entitled **Student Notes**. If you have time, read through the Student Notes before each lecture. This will help you get the most out of your classroom time. After lecture, use the Student Notes to reinforce and review what you have learned.

# Module IN — Introduction

HP supports several families of HP-UX products. HP-UX runs on the HP 9000 Series 300/400, HP 9000 Series 700, and HP 9000 Series 600/800 systems. These families share much in common. As such, there is much overlap in the knowledge required by a system administer of any of these families. The course you are attending, either HP-UX System Administration for the HP9000 Series 600/800 or HP-UX System Administration for the HP9000 Series 300/400/700, will have material in common with its complement course for the other system family along with modules written expressly for your particular series of computer. Every effort has been made to clearly mark any material that is *not* applicable to all families.

## IN-2. SLIDE: The Role of the System Administrator

2



51436 IN-2. 2 

## Student Notes

The system administrator is responsible for setting up and maintaining the system. Not only must the administrator understand both hardware and software, but he/she must also understand the needs of the user community.

Since many of the tasks associated with these responsibilities require access to commands that should not be available to every one, the system administrator needs special access to the system. This access is called superuser or root access.



10

## IN-3. SLIDE: Hardware Responsibilities

3

### Hardware Responsibilities

- Verify that system hardware is installed correctly and tested
- Verify that peripherals are installed correctly and tested
- Monitor performance of hardware components
- Arrange for repair in event of hardware failure

51436 IN-3. 3 

## Student Notes

The system administrator of an HP-UX system is responsible for verifying that the HP 9000 system **hardware** has been installed correctly and thoroughly tested. This must be done before an HP-UX operating system can be loaded.

The administrator may not be the person who actually installs the hardware. Often a Hewlett-Packard Customer Engineer will perform the installation of the hardware and possibly of the system software. Once the system is operational, the administrator must monitor the performance of the various hardware components. If a hardware failure occurs, the administrator should attempt to isolate the problem as much as possible. Depending upon the service agreements in place, the administrator may schedule a customer engineer to make necessary repairs.

The system administrator must know some basic things about the system hardware to be effective in his job. We will cover the hardware at an overview level in this course.

## IN-4. SLIDE: Software Responsibilities

4

### Software Responsibilities

- Install the HP-UX operating system.
- Configure the HP-UX operating system.
- Create file systems.
- Detect and correct file system errors.
- Monitor file system usage and growth.
- Design and implement backup and recovery routines.
- Configure and maintain printer spooler software.
- Install and maintain network communication software.
- Update the HP-UX operating system for new releases.
- Install and update application software.

51436 IN-4. 4 

## Student Notes

One of the primary responsibilities of the administrator is to install the HP-UX operating system software. The operating system is supplied on a distribution medium that can be one of many forms. The software is distributed on a series of magnetic tapes or on optical CD- ROM discs. The distribution media may be:

- Compact Disc - Read Only Memory (CD-ROM) disc
- Digital Data Storage (DDS) tapes based on Digital Audio Tape (DAT) technology
- 1/4 inch cartridge tape
- 1/2 inch mag tape

The type of media used varies based on the system type, we will discuss this in detail later. The system software must be installed onto a hard disk using a special command set by the system administrator.

Once the HP-UX system has been installed, it will be necessary to configure the software by altering certain values or parameters. This is done to allow the operating system (HP-UX) to recognize additional devices or to increase the efficiency of the system. (Note that an in-depth discussion of optimizing the kernel by modifying its tunable parameters is outside the scope of this course and will not be covered.)

We'll see later that HP-UX utilizes a data organization scheme called a file system. Since the file system is where all of the system and user data is stored, it is quite important to ensure that the integrity of the file system is maintained. The administrator should implement procedures which will detect any errors or corruption in the file system. If problems are found, the administrator must ensure that corrective action is taken.

A file system is of a finite size and resides usually on a locally connected disk drive. The available space in a file system must be monitored by the system administrator. Procedures should be employed to archive and/or remove obsolete and unused files so the amount of free space is not compromised. Files that tend to grow in size should be observed as well.

As in any computer environment, accidents will happen. Files may be removed accidentally, file systems may become corrupted, hardware may fail. It is very likely that any of these cases may result in a loss of data. Consequently, it is the administrator's job to ensure that adequate backups are maintained. Then, if there is a loss of data, recovery procedures can be employed.

Most HP-UX systems include peripheral devices for hard copy output such as line printers and laser printers. The system administrator must manage the software which sends output to the printing devices.

Different HP-UX systems communicate with one another across an electronic communications mechanism called a network. The network allows electronic mail and files to be transmitted from machine to machine. Networking software must be installed, configured, and monitored by a system administrator.

Usually once a year, and occasionally several times a year, Hewlett Packard will release an update to the HP-UX operating system software, and to many of the subsystems and applications programs.. An update may enhance or modify existing system features or add new capabilities. The system administrator is responsible for installing each software update so the HP-UX system available to the user community contains the latest version of the software.

The fact that you receive a new HP-UX release does not automatically mean you must install it right away. It is your responsibility to ensure that interactions and dependencies between applications and system software are not detrimental. A new HP-UX release is often followed several months later by releases for applications that make them compatible with the new HP-UX release. Installing the new HP-UX release too soon may break an application that worked fine on an older release. Obtaining and sharing this information is your responsibility.

## IN-5. SLIDE: Responsibilities to The Users

5



### Student Notes

Once the HP-UX system has been installed, certain modifications are required to allow a user to access the system. These modifications must be performed by the system administrator.

The administrator must, to the greatest extent possible, tailor the system to the needs of the user community. The system administrator should analyze the intended use of the system, and should be aware of the number of users on the system, the characteristics of each user, the system resources and peripherals required by each user, and the data/programs that must be shared by various user groups.

As system administrator you will be looked upon as the resident HP-UX expert. Many users will assume that you know everything about the system and will view you as a "guru." This occurs independently of

your knowledge level. To many users, the fact that you have been trained, in their minds, means that "you must know" more than they do.

You may be asked many questions such as, "How do I do ... ?" and receive comments such as "My terminal is broken" and "I forgot my password." The problem solving and consulting aspects of the administrator's role can be the most challenging as well as the most enjoyable parts of the job. An important message to convey early on is that the HP-UX reference manuals are available on-line on an HP-UX system. Get the users used to at least trying to "look it up for themselves" before coming to you.

Since you will implement the policies and procedures of your company/organization, be aware that these policies and procedures take precedence over the things that HP will tell you in this class. We can only recommend certain administration practices. If our recommendations are in conflict with your company practices, clearly you have to follow your company guidelines.

## IN-6. SLIDE: Typical Course Flow

6

### Typical Course Flow

Spool Disk

- Hardware Overview
- Boot Up
- Installing HP-UX
- Device Files
- HP-UX Startup Sequence
- System Shutdown
- File System Structure
- File System Creation
- Maintaining the Integrity of a File System
- Backing Up the System
- Post-Installation Procedures
- Memory Management
- Swap
- Kernel Reconfiguration
- Updating HP-UX
- Common Administrative Tasks
- Managing the LP Spooler

51436 IN-6. 6 

## Student Notes

This course is designed to teach you the basic concepts and tasks needed to administer an HP-UX system. By no means do we cover everything there is to know about system administration, but the skills you learn in this class will get you going and give you a solid base upon which to build more knowledge.

The course is divided into the following modules:

- Hardware Overview
- Boot Up
- Installing HP-UX
- Device Files

# Module IN — Introduction

- HP-UX Startup Sequence
- System Shutdown
- File System Structure
- File System Creation
- Maintaining the Integrity of a File System
- Backing Up the System
- Post-Installation Procedures
- Memory Management
- Swap
- Kernel Reconfiguration
- Updating HP-UX
- Common Administrative Tasks
- Managing the LP Spooler
- Optional—Setting Up a LAN
- Optional—How to Administer UUCP
- Optional—Creating a Cluster
- Optional—System Accounting



14

## IN-7. SLIDE: HP-UX Manuals For System Administrators

7

### HP-UX Manuals For System Administrators

Common to all HP-UX systems:

- *HP-UX Reference sections 1M*
- *How HP-UX Works, Concepts For the System Administrator*
- *Managing Clusters of HP 9000 Computers, Sharing the HP-UX File System*
- *HP-UX System Security*
- *Solving HP-UX Problems*
- *Error Message Catalog*
- *A Beginner's Guide to HP-UX*

Separate versions for Series 300/400, 700 and 600/800:

- *System Administration Tasks Manual*
- *Installing Peripherals*
- *Installing and Updating HP-UX*
- *Master Index*
- *Finding Information*

51436 IN-7. 7 

## Student Notes

Often during the administration of your system, you will need to reference the documentation. There are many HP-UX manuals. How do you know which manual to check? The slide lists manuals specifically targeted towards administrators. You should be familiar with these manuals and their content as they can be of tremendous help to you.

Note that some of the HP-UX manuals are specific for one computer family or architecture. When consulting the manuals be sure that you are using the manual designed for your computer family. Other manuals that cover more general concepts are provided by HP as a common HP-UX manual and the contents apply to all members of the HP 9000 HP-UX computer family.

### HP-UX Reference

This manual set (two volumes) is intended for all HP-UX systems. It is intended as reference material and is most useful to experienced users. It is not designed to serve as a learning tool for beginners. The manual set is comprised primarily of HP-UX command descriptions and supporting information. Remember that the content of this manual set is also available on-line by using the `man(1)` command.

The HP-UX Reference Manual is divided into eight sections:

- Section 1: User Commands
- Section 1M: System Administration Commands
- Section 2: System Calls
- Section 3: Subroutines
- Section 4: File Formats
- Section 5: Miscellaneous Facilities
- Section 7: Device Files
- Section 9: Glossary

This course will draw heavily upon section 1M, System Administration Commands. Section 1M contains information on those commands that are used primarily by a system administrator. Section 4, File Formats, is also invaluable to the system administrator as it contains information on most of the configuration files that you will be responsible for maintaining.

### How HP-UX Works, Concepts for the System Administrator

This manual explains HP-UX system concepts useful to programmer's, users and system administrators. Concepts covered include system startup, system shutdown, login, processes, run- levels, memory management, file system, device files, using peripherals, networking, clusters, and system accounting.

### Managing Clusters of HP 9000 Computers, Sharing the HP-UX File System

This manual is essential for understanding how to manage a diskless cluster environment regardless of which HP-UX system you administer. However, an administrator not in a cluster, also called a standalone system, will not find this information applicable to his situation.

### HP-UX System Security

*HP-UX System Security* is a reference book of procedures and guidelines essential to administering a secure system. The manual provides detailed information on implementing HP's full range of security features, such as auditing and refined control mechanisms for file access.

### Solving HP-UX Problems

This manual is a troubleshooting manual for the HP-UX system Administrator.

### Error Message Catalog

This common manual provides a complete list of all system error messages you are likely to encounter in HP-UX.



15

# Module IN — Introduction

## A Beginner's Guide to HP-UX

This manual is for beginners on HP-UX, but you should be thoroughly familiar with ts contents and examples, since it may be the source of questions by novice users. This manual combines previous documents for beginners guides on text editing, getting started with HP-UX commands, beginners guide to shell, the "vi" quick reference, and much of the content of "Introducing UNIX System V".

## System Administration Tasks Manual

This manual cover most system administration tasks in detail. Topics covered include constructing and customizing an HP-UX system, updating HP- UX, starting and stopping HP-UX, managing run-levels, managing groups, users, file systems, and clusters, and setting up and administering backups and the LP spooler.

## Installing Peripherals

This manual provide step-by-step instructions on how to configure peripherals, such as terminals, printers, plotters, as well as disk and tape drives.

## Installing and Updating HP-UX

This manual covers what its title states: how to install, update and modify the HP-UX operating system.

## Master Index

This manual provides a complete index of the contents of all the relevant HP-UX manuals. When used with finding Information, you should be able to identify the manual and the chapter where a particular concept is discussed. It is in effect a multi-volume index of relevant HP-UX manuals for a particular family of HP-UX systems.

## Finding Information

This manual is designed to help you locate information that will help you work with the HP- UX operating system. There are many HP-UX manuals available, and you may find it difficult to decide which ones to use. This manual lists all of the HP-UX documentation available with the part number, intended audience, and tasks and concepts covered. If you look through this manual, you should be able to figure out what other manuals you need.

16

# Module HW — Hardware (Series 300/400/700)

## Objectives

Upon completion of this module, the student will be able to do the following:

- Identify the major hardware components of HP 9000 Series 300, 400, and 700 computers.
- Define hardware addressing terms such as:

  Series 700 - SBM and Function number
  Series 300/400 - Select code
  All systems - Bus address

- Change bus addresses.
- Install interface cards while observing electrostatic precautions.
- List peripherals frequently found on HP 9000 systems.

# Module HW — Hardware (Series 300/400/700)

## Objectives

Upon completion of this module, the student will be able to do the following:

- Identify the major hardware components of HP 9000 Series 300, 400, and 700 computers.
- Define hardware addressing terms such as:

  Series 700 - SBM and Function number
  Series 300/400 - Select code
  All systems - Bus address

- Change bus addresses.
- Install interface cards while observing electrostatic precautions.
- List peripherals frequently found on HP 9000 systems.

## HW-1. SLIDE: Series 700 System Processors

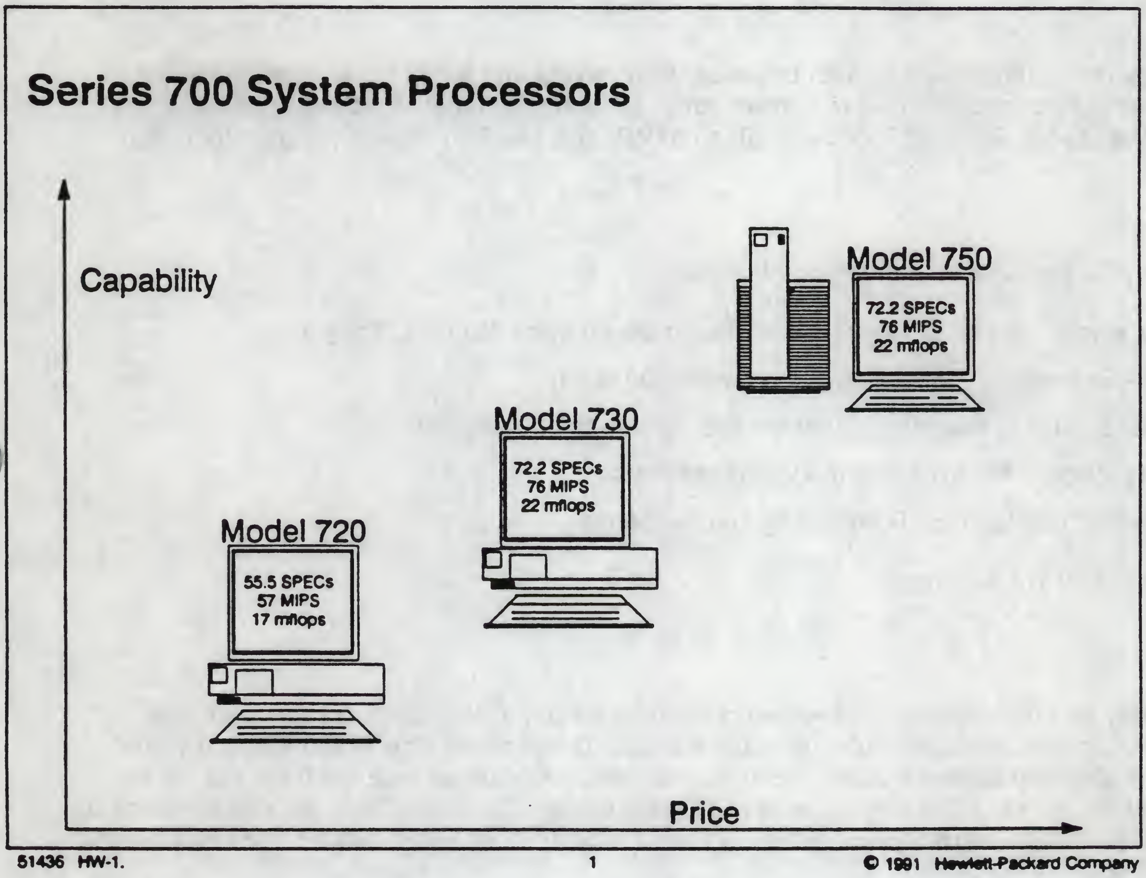


## Student Notes

The Series 700 includes these fundamental hardware components:

- System Processing Unit (SPU)
- Memory
- Console
- Disk (optional)

### System Processing Unit

In 1991, Hewlett-Packard announced a break through in price- performance, the HP 9000 Series 700 families of SPUs. The Series 700 is based on the second generation of HP's PA-RISC processor, PA-RISC1.1 The original PA-RISC was introduced in 1985 with the HP 9000 Model 840. The 50 and 66MHz frequency of the Series 700 is a result of a low-power CMOS technology and advanced package design methodologies.



17

The Model 720 runs at 50 MHz while the Model 730 and Model 750 run at 66 MHz. The Models 720 and 730 are desktop implementations with reduced I/O expandability versus the Model 750 which is a deskside model with four slots available for industry standard EISA plug in boards.

### Memory

Memory on the Series 700 is implemented with two-way interleaving and a 128-bit wide memory bus. ECC error detection and correction is used to correct single bit errors and detect all multiple bit errors. RAM capacities for the Models 720 and 730 are 16MB to 64MB, and 16MB to 192MB for the Model 750.

### Console

A system console may be terminal or a bit-mapped display.

The Series 700 offers a wide choice of graphics options for bit-mapped displays. They are:

| | |
|---|---|
| GRX | High-performance grayscale display with 256 shades |
| CRX | 8-plane color display for 2D vector and 3D wireframe graphics |
| PVRX | Entry display for 3D solids modelling graphics |
| TVRX | High-end display for 3D modelling and rendering |

All of these displays are 72 Hz monitors.

### Disk

Series 700 systems may be configured as diskless workstations making disk support optional. When configured for a disk, support is via the industry standard SCSI-II interface. The Model 720 and Model 730 can each accommodate up to two 3 1/2"internal SCSI-II disks. Additional support for up to seven 5 1/4" disks is provided externally by the optional SCSI-II EISA board. The Model 750 can accommodate up two 5 1/4"internal SCSI-II disks with a possible four optional SCSI-II EISA boards which each can support up to seven 5 1/4" disks.

### Measurements of Performance

While the only true test of performance is the application that the user will actually run, porting of a custom application to several hardware platforms in order to calculate performance numbers may prove impractical. To aid the user in his buying decision, the Computer Industry over the years has developed several units of measures to compare relative performance of hardware.

| | |
|---|---|
| MIPS | Usually MIPS (Millions of Instructions Per Second) are not an actual measure of instructions per second, but a number relative to the performance of a DEC VAX 11/780 which is considered by the industry to be a 1 MIP machine. These particular MIPS ratings are based on Dhrystone MIPS relative to a VAX 11/780. |
| MFLOPS | These Megaflops or MFLOPS (Millions of Floating Point Operations Per Second) measurements are based on the Linpack Double Precision Floating Point benchmark numbers. |

SPECs

SPECmarks are a suite of benchmarks developed by Systems Performance Evaluation Cooperative (SPEC) which is non-profit corporation whose members are primarily computer system manufacturers.



16

## HW-2. SLIDE: Series 400 System Processors

8

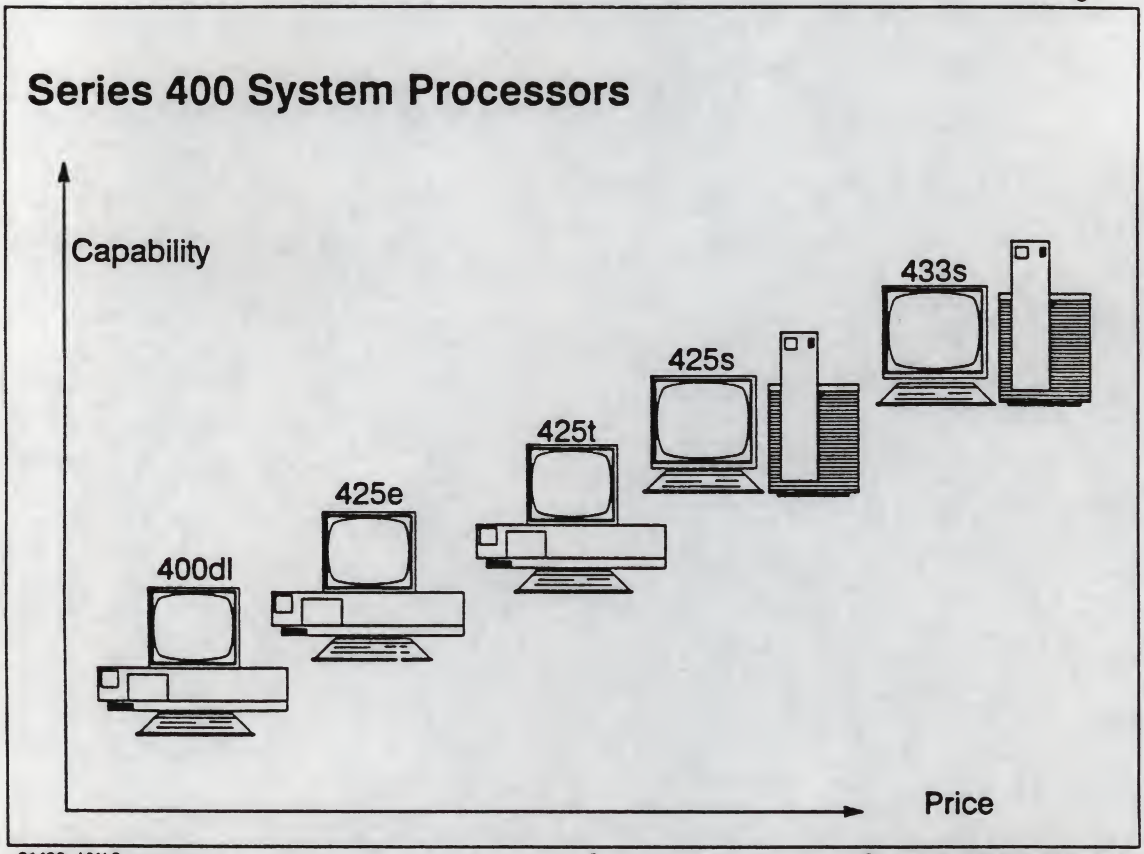


51436 HW-2. 2 

## Student Notes

Series 400 Model 400dl is an aggressively priced entry-level 32-bit diskless workstation. The Model 400dl is a pre-configured system for monochrome graphics and X Window System support. Like more expensive models, the 400dl features a MC68030 50MHz processor and a floating point coprocessor. The 400dl is not expandable to a standalone configuration or to support a color display.

In the center of the chart are mid-level performing Series 400 systems. The Models 425e, 425t, and 425s have use a MC68040 25MHz processor. These models offer greater expandability than the 400dl. The desktop models allow a smaller low profile package, while the deskside models allow more IO expansion and larger internal disks.

On the right is the higher performing Model 433s which uses a MC68040 33MHz processor. The 433s system also provides greater possibilities for system expansion compared to the 400dl, 425e, and 425t desktop models.

Desktop models, 423e, and 425t, support advanced performance levels, internal disk drives, additional I/O via SCSI adapters, and a wide range of optional graphics accelerators for interactive manipulation of 3D images and solid models.

Deskside (suffix "s") models such as the 400s and 433s support advanced performance and graphics products for interactive manipulation of 3D images and solid models.

The model 400dl maintains its position as the low cost entry level Series 300/400 workstation while the Model 425t provides the best price/performance value in the Series 300/400 family.

**Table HW-1.**

| System | CPU | clock | MIPS | SPECS | MFLOPS |
|---|---|---|---|---|---|
| 400dl | 68030 | 50 MHz | 12 | 4.1 | 0.5 |
| 425e/t/s | 68040 | 25 MHz | 20 | 11.0 | 2.7 |
| 433s | 68040 | 33 MHz | 26 | 15 | 3.8 |

## HW-3. SLIDE: Series 300 System Processors

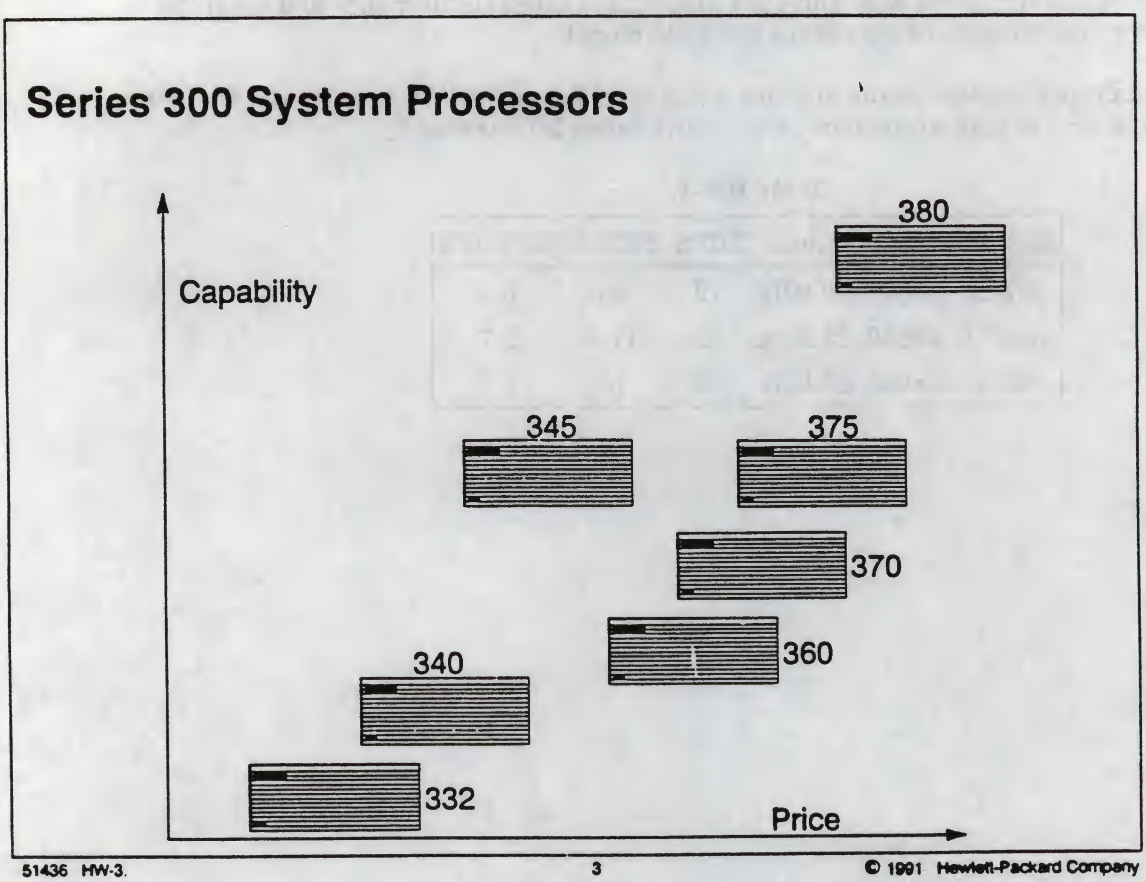


## Student Notes

The slide shows the several representative Series 300 workstation computers. On the left side of the chart are entry level workstations with 3 MIPS performance levels. Model 332 and R/332 are pre-configured systems for 2D graphics and instrument controller applications. These are low cost, entry level models which feature MC68030 processors with optional floating point coprocessors.

Systems based on Models 330, 340 and 360 offer mid-level performance within the Series 300 family, which is 3 to 5 MIPS performance. They also provide greater possibilities for system expansion of memory and I/O compared to the model 332. The mid level models offer higher-resolution graphic displays, and optional accelerators for faster computation and higher graphics performance than the entry-level workstations.

The 345, 370 and 375 models also feature the MC68030 processor. The model 370 processor operates at 33 megahertz with floating point coprocessing for 8 MIPS performance. The Model 345 and Model 375 contain a 50- megahertz MC68030 CPU and 50-megahertz MC68882 Floating point coprocessor and an

on-board 32-KB processor cache. These models function at the 12 MIPS and 0.5 MFLOPS performance level.

The model 345 is high-performance, low-cost desktop Series 300 workstation, which supports an optional 200-MB built-in disk. The Model 375 is housed in an expandable, flexible and rack mountable package while the model 345 offers the same performance in a small desktop cabinet without expansion capability. The model 375 supports advanced performance and graphics products for interactive manipulation of 3D images and solid models.

A daughter board for the main CPU board in the model 375 contains a 50-megahertz MC68030, a 50 megahertz MC68882 and an on-board 32-KB processor cache. The hardware integration using ASICs (Application Specific very large scale Integrated Circuits) reduces cost and space and improves reliability. An extra DIO-II slot is available in the Model 375, which is another advantage over the Model 370.

A plug-in CPU upgrade for the Model 375 is available. The daughter board for the main CPU board is replaced by a 25- megahertz 68040 processor to raise the CPU to 20 MIPS performance levels. The upgrade replaces the on board cache and the floating point coprocessor, which are both integrated into the 68040 processor. The upgraded 375 system is the model 380 system.The upgraded floating point performance is 2.7 MFLOPS.

All current Series 300 system processors are based on the 68030 or 68040 processor. Earlier versions of Series 300 systems utilized either 68010 or 68020 CPUs.



20

## HW-4. SLIDE: View of Models 720 and 730

### View of Models 720 and 730

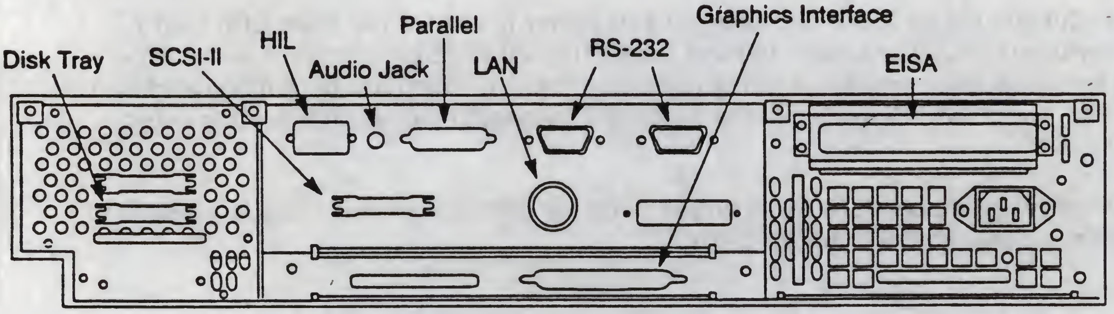


Model 730 and 720 Rear Panel

51436 HW-4. 4 

## Student Notes

The Models 720 and 730 are desktop workstations that consist of the following hardware:

- internal storage
- memory slots
- core I/O board
- graphics board
- EISA slot
- power supply

The Models 730 and 720 have space for two 3 1/2 inch internal storage devices in single slide out tray that can easily be locked up for secure data applications. At introduction, the internal storage in the back of the tray is either a 210MB disk or 420MB disk. The front can contain a 210MB disk, 420MB disk,

or a 3 1/2 inch floppy disk drive. These internal storage devices are single-ended SCSI II. Shipped from the factory, the system disk is at bus address 6.

Memory cards may occupy any slot pair of the four pairs of slots.

The core I/O board supports HIL, audio, two asynchronous RS-232C ports, an IEEE 802.3 Ethernet LAN interface configured with ThinLAN, one Centronics parallel port, and an external SCSI II port. The SCSI II port is also connected to the internal storage device(s).

The graphics board at introduction supports HP's GRX grayscale, CRX color 2D/3D, PVRX color 3D, and high-performance TVRX color 3D monitors.

The EISA adapter supports one EISA interface card on the Model 730. This is optional on the Model 720. Cards currently supported at introduction include EISA differential SCSI II Host Adapter, EISA HP-IB Host Adapter, and EISA LAN/9000.

The power supply is auto selecting between 115V and 230V.

## Models 720 and 730 System Options Diagram

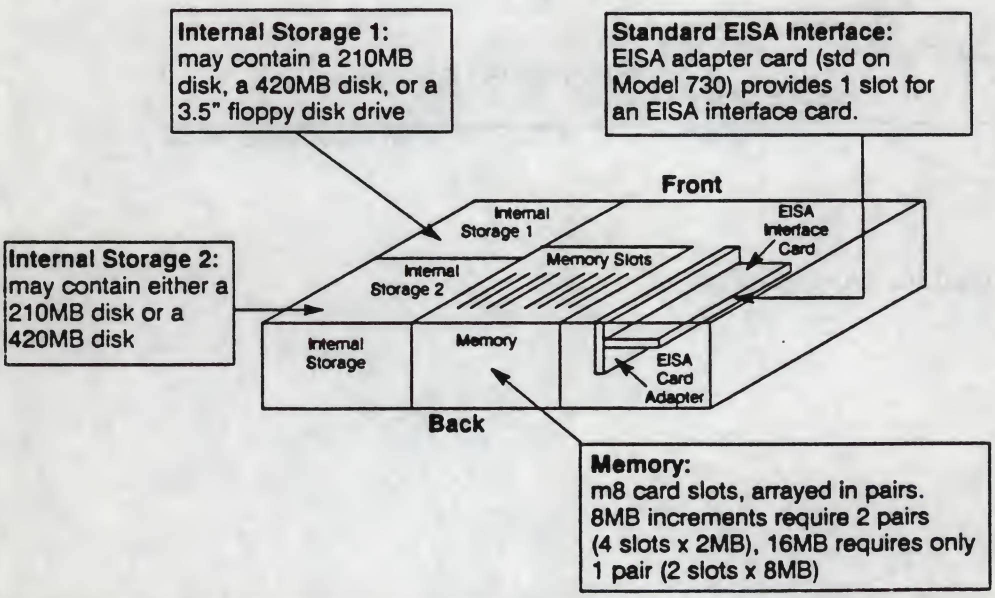


51436 HW-5. 5 



21

## HW-5. SLIDE: View of a Model 750

### View of Model 750

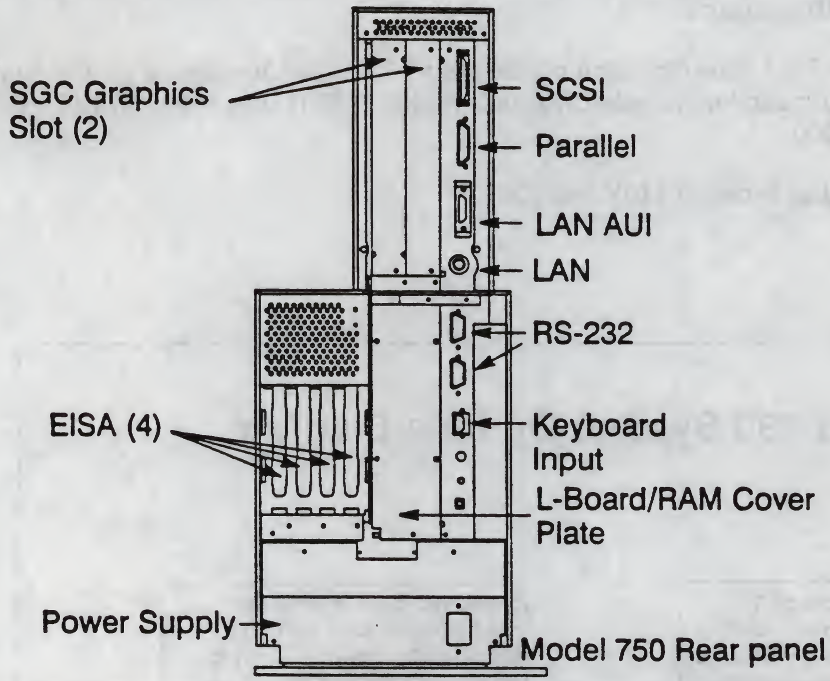


51436 HW-6. 6 

## Student Notes

The Model 750 deskside workstation consists of:

- removable media
- internal storage
- memory slots
- core I/O board
- graphics board
- EISA slots
- power supply

The Model 750 has space for removable media. This space may contain a single DAT/DDS Format tape drive or a CD-ROM and a floppy disk drive.

The footprint for an internal storage device on the Model 750 is larger than the Models 730 and 720. These two storage bays accommodate two 5 1/4 inch internal SCSI-II disks. Disks included in this 5 1/4 inch formfactor category include 660MB disk, 1.3GB disk or DAT drive.

Memory cards may occupy any slot pair of the six pairs of slots.

The core I/O board supports the same connections as the Models 730 and 720.

Graphics support is also the same as the Models 730 and 720.

The ESIA adapter on the Model 750 supports four EISA slots. The same EISA cards are supported on the Model 750 as the Models 730 and 720.

## Model 750 System Options Diagram

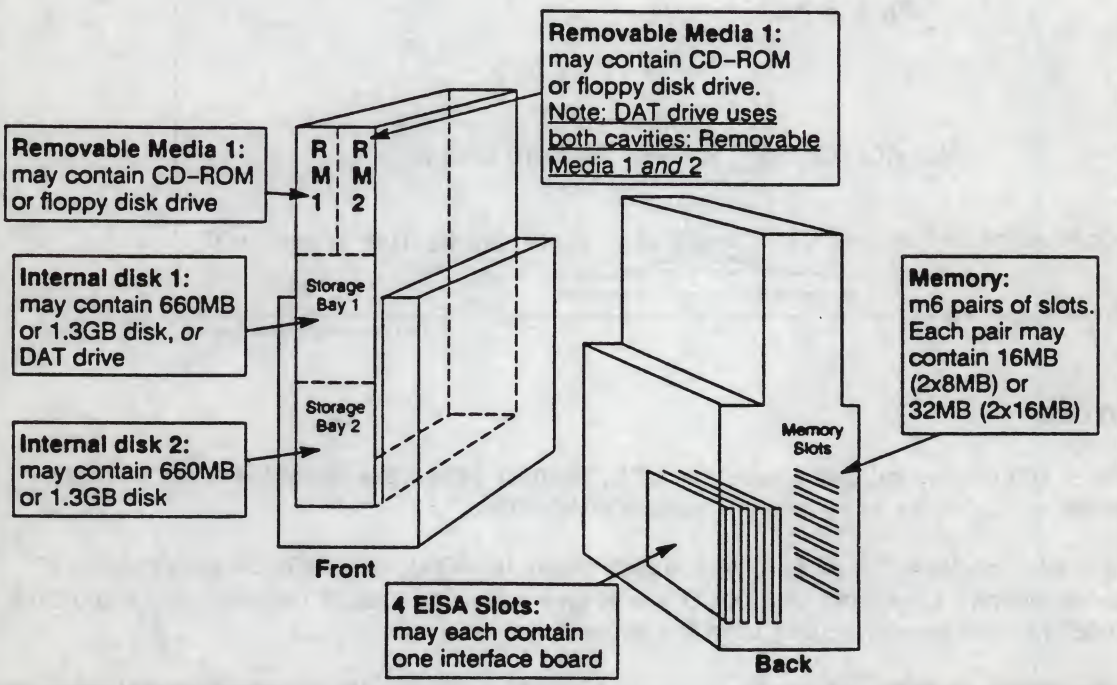


51436 HW-7. 7 



22

## HW-6. SLIDE: View of Models 400dl and 425t

9

### View for Models 400dl and 425t

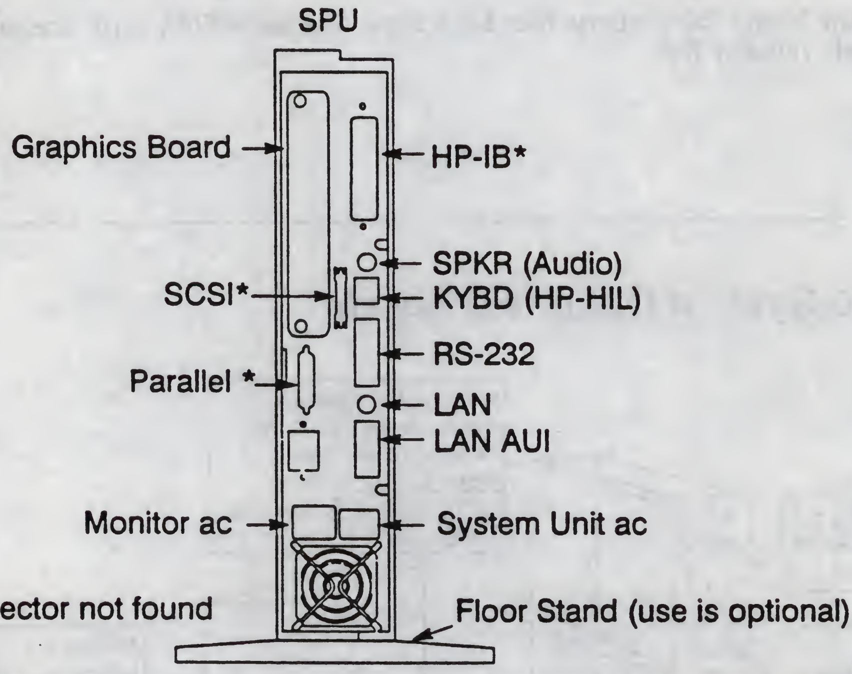


* Interface connector not found on 400dl

51436 HW-8. 8 

## Student Notes

The HP 9000 Series 400 system cabinet houses the CPU, Memory Management Unit, system memory (RAM), and in some cases one or more internal system disk drives.

This is an example of a model 400dl or 425t back panel. Some HP 9000 Series 400 desktop (suffix "t" and suffix "dl") models may vary from this, but this will give you some idea of the connectors and their location. The 400dl has fewer connections than are shown here.

The HP 9000 Series 400 system processors have internal interface circuits that support several types of interfaces, and the corresponding interface connectors appear on the back panel. The interface circuitry and connectors of a Series 400 system is not on a removable system card as we saw in the Series 300 SPUs.

The interface connectors are listed in clockwise order starting at the top right.

- Standard speed HP-IB (optional)*

- SPKR connector (audio output output)
- KYBD connector (HP-HIL Human Interface Link)
- RS-232 connector (serial data communications)
- LAN connector (Local Area Network Thin LAN type BNC)
- LAN AUI connector (Local Area Network Adapter Unit Interface)
- AC power in (system input power connector)
- AC power out (power connector for monitor)
- HP Parallel interface connector (Centronics compatible)*
- DIO-II interface board slot (usually a graphics adapter)
- SCSI interface for external digital data storage.*

Items marked with a "*" are not found on 400dl

The 400dl and 425t models are non-expandable models designed for table top use. They may be used with the deskside cradle supplied with the SPU. See the new section on deskside systems for a better view of Series 400 family expansion capabilities.

## HW-7. SLIDE: View of Model 425e

10

### View for Model 425e

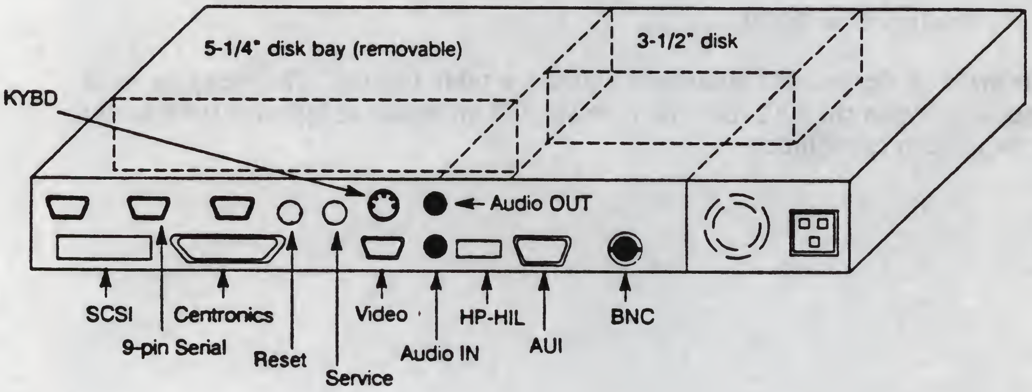


51436 HW-9. 9 

## Student Notes

At its introduction, the Model 425e was a completely new workstation design based on the Motorola 68040 25MHz processor for a low price-point entry.

This model includes the following standard interfaces:

- Three 9-pin serial ports.
- PC keyboard
- Audio in/out
- SCSI
- Centronics
- Video
- HP-HIL
- LAN AUI connector for MAU
- LAN BNC connector for ThinLAN

## HW-8. SLIDE: View of Models 425s. and 433s - DIO-II Card Cage

11

### View of Models 400s, 425s, and 433s DIO-II Card Cage

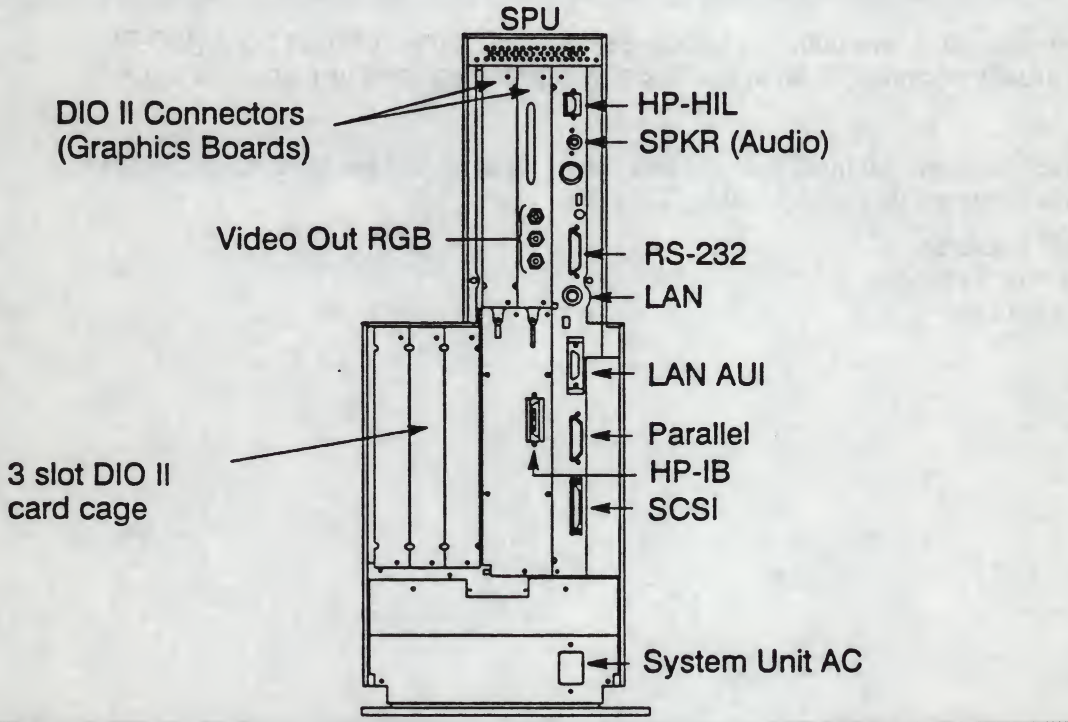


51436 HW-10. 10 

### Student Notes

The System cabinet houses the CPU, Memory Management Unit, system memory (RAM), and in some cases one or more internal system disk drives or other digital data storage devices.

The internal system interface circuits support a number of interfaces, and a number of interface connectors appear on the back panel. Unlike the Series 300, the interface logic is not on a removable system card.

The interface connectors are listed in clockwise order starting at the top right.

- HIL connector (HP-HIL Human Interface Link)
- SPKR connector (audio output output)
- Apollo Domain KYBD (Not used in HP-UX)
- Service Mode switch (Not used in HP-UX)
- Reset switch (Not used in HP-UX)
- RS-232 connector (serial data communications)



24

- LAN connector (Local Area Network Thin LAN type BNC)
- LAN AUI connector (Local Area Network Adapter Unit Interface)
- Parallel interface connector (Centronics compatible)
- SCSI interface connector.
- AC power (system input connector)
- HP-IB connector (Optional)
- Two DIO-II bus slots.

All Series 400 deskside systems have two built-in expansion slots which support Direct I/O II (DIO-II) cards. One DIO-II slot is usually occupied by an optional graphics processor card or a graphics adapter card.

The 425s and 433s SPUs are deskside cabinets that can have their basic backplane capacity expanded via several different optional expanders that are available. They are

3-slot INTERNAL DIO-II Expander
4-slot INTERNAL EISA Bus Expander
DIO-II to DIO adapter card cage

## HW-9. SLIDE: View of Models 400s, 425s, and 433s - EISA Card Cage

12

### View of Models 400s, 425s, and 433s EISA Card Cage

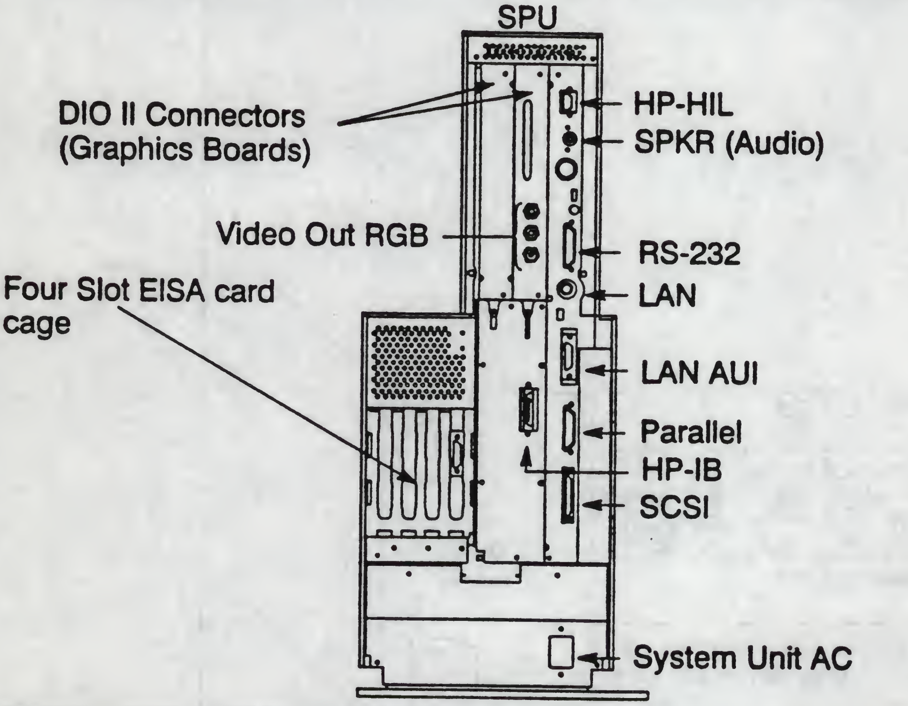


51436 HW-11. 1 

## Student Notes

This is an example of a model 433s back panel with an optional four slot EISA card cage.

The interface connectors are unchanged with respect to the previous slide except for the: Four optional EISA interface board slots in the main SPU cabinet/enclosure.

Only the 400s and 433s models support the optional EISA card cage.

The 4-slot INTERNAL EISA Bus Expander is an optional expander that provides four EISA slots. This hardware configuration is only supported on the HP 9000 Series 400s and 433s systems under HP-UX at release 8.0.



215

## HW-10. SLIDE: View of a Model 360

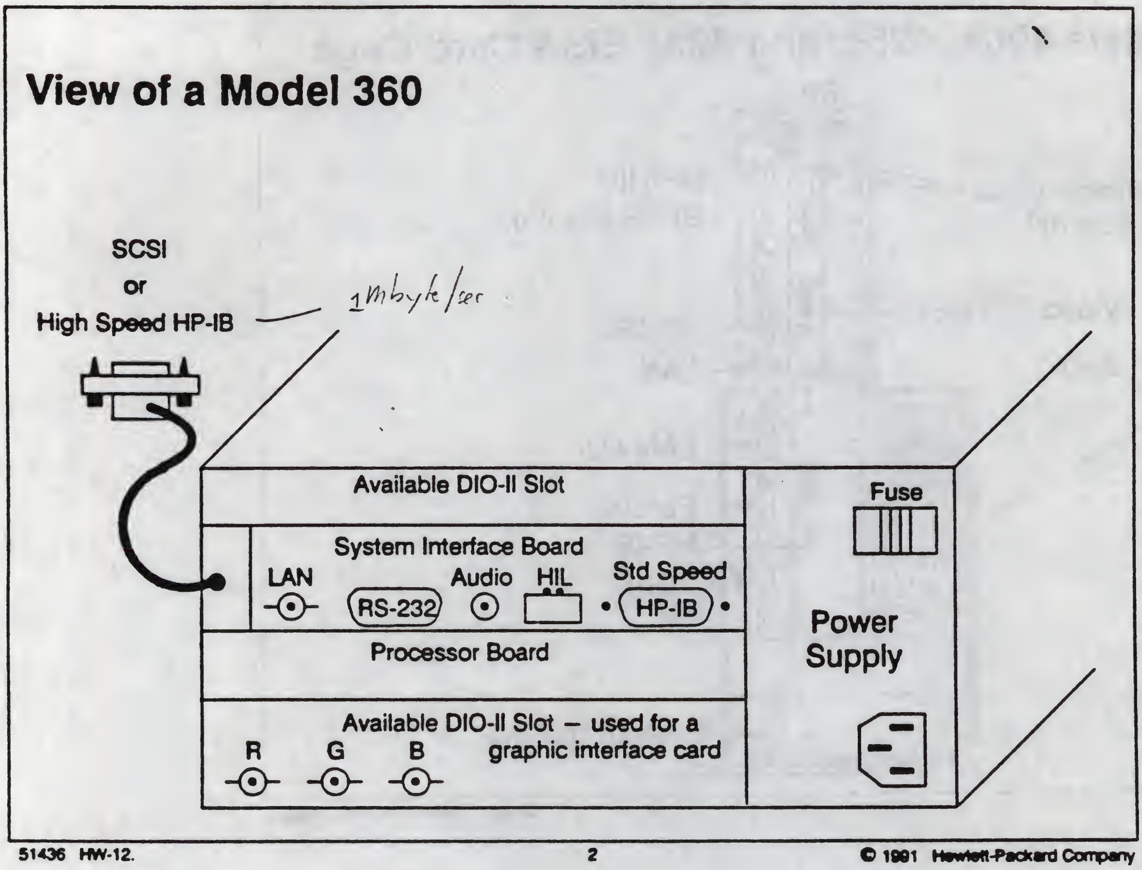


## Student Notes

In the HP 9000 Series 300 computer family, the processor board houses the CPU, Memory Management Unit, and, in some cases, the system main memory (RAM).

This is an example of a Model 360 backplane. Other Series 300 models vary somewhat from this, but this will give you a good idea of the typical back panel connectors and the various cards that may be plugged into the HP 9000 Series 300 backplane. The Model 360 System Interface Board supports a number of interface connectors:

High Speed disk interface (SCSI or HP-IB)
LAN (Local Area Network)
RS-232 Serial Data Communications
Audio output
HIL (Human Interface Link)
HP-IB (HP Interface Bus)

One significant difference is that the 345 and 375 models have a SCSI low density connector and a parallel port on the back panel.

BUSES (sometimes called the system backplane) are needed to transfer data between the System Processor card, the main memory (RAM) cards, and the input/output (I/O) cards. There are two different buses in the HP 9000 Series 300 family: Direct I/O (DIO) bus and Direct I/O II (DIO-II) bus.

The DIO Bus is a 16-bit bus and is used by most Series 300 interfaces and some accessory cards. The DIO bus cards are physically smaller than the DIO-II cards in both width and depth. DIO Bus transfers are slower than DIO-II transfers.

The DIO-II bus is used in the models 330, 340, 350, 360, and 370. It is a 32-bit, double width bus and is used by floating point math and graphics accelerator cards, system cards (such as memory expansion cards in some systems), and can support an eight port DIO-II MUX. DIO-II Bus transfers are faster than DIO transfers.

The number of available DIO-II slots will vary depending upon the Series 300 model. Some models do not support any available slots (models 340 and 345 are examples). If present on your SPU, an available DIO-II slots may be used for:

A supported Graphics Interface Adapter
Additional Memory Expansion (only on some system units)
HP-98638A 8-Port RS-232C/RS422 Serial Multiplexer
HP-98248A Floating Point Accelerator Accessory Card
HP-98242A DIO-II to DIO adapter card cage

## HW-11. SLIDE: Series 700 Bus Architecture

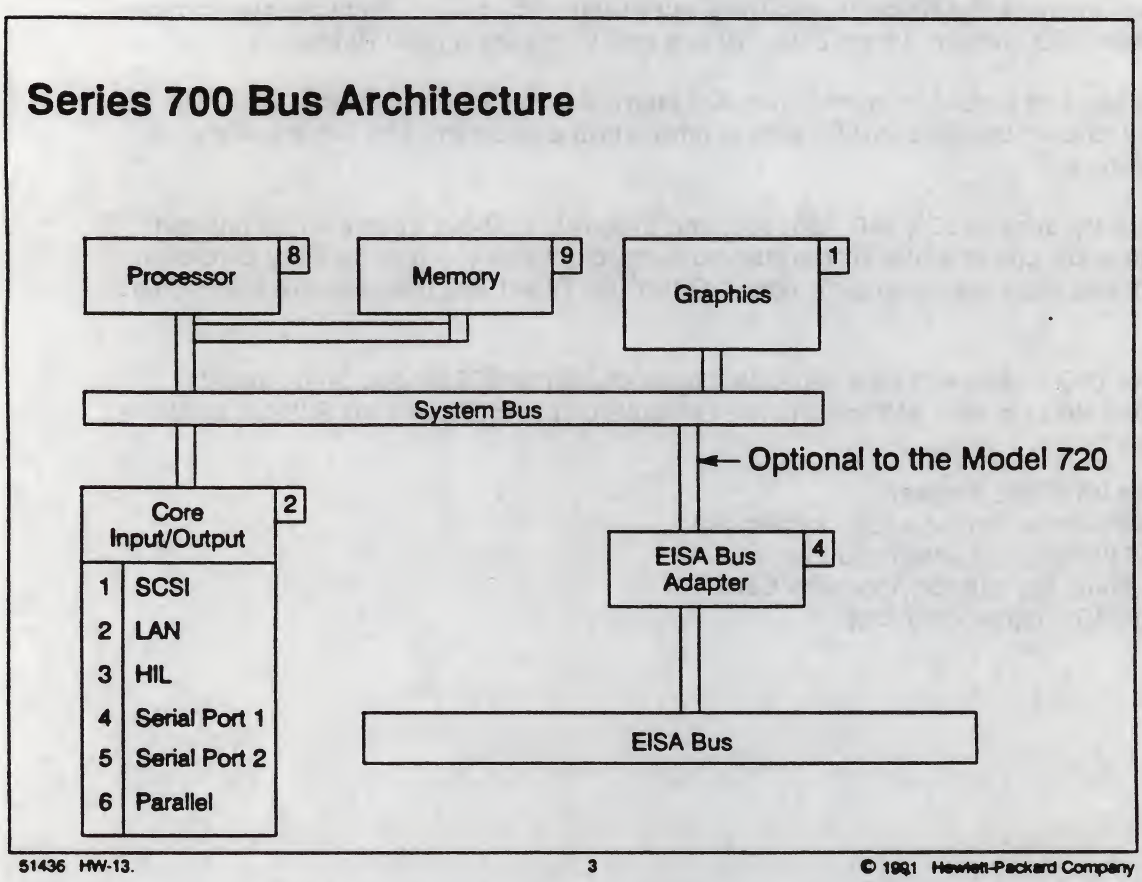


### Student Notes

PA-RISC I/O Architecture is based on functional entities called **modules.** A minimal system consists of a processor, memory, and I/O modules attached to a single system bus. Specifically on the Series 700, the major modules attached to the system bus are the processor, memory, graphics, EISA Bus Adapter and the core Input/Output system. In the figure the small boxes appended to the modules are the System Bus Module (SBM) numbers. SBM numbers are used to address the hardware attached to the system bus.

Graphics, processor, and memory are at SBM numbers 1, 8, and 9 respectively. The core I/O at SBM number 2 contains separate interfaces, each of which are represented by a **function** number. The EISA Bus Adapter at SBM 4 is responsible for conversion of EISA protocol to and from PA-RISC I/O protocol. The EISA protocol is an industry standard protocol used by the EISA cards plugged into slots of the EISA bus. In the case of an EISA card, the slot number is also required for hardware addressing.

## HW-12. SLIDE: Default Select Codes for the Series 400 and 300

13

### Default Select Codes for the Series 300 and 400

| Interface | Select Code | |
|---|---|---|
| High-speed HP-IB disk | 14 | 0x0E |
| SCSI interface | 14 | 0x0E |
| Standard-speed HP-IB | 7 | 0x07 |
| LAN | 21 | 0x15 |
| HP-HIL | 0 | 0x00 |
| RS-232 (built-in) | 9 | 0x09 |
| Centronics | 12 | 0x0C |

11 x16⁰ x16¹

51436 HW-14. 4 

## Student Notes

The Series 400 and 300 locates interfaces by the interface's **select code**. Each interface has its own unique select code.

Select codes are needed for the built-in interfaces in the SPUs and for optional add-in interface cards. Some select codes are preset by the factory as indicated in the slide. It is recommended that these select codes be changed only if the need presents itself. For instance, you may have more than one LAN card; they cannot both be at select code 21.

Depending on the processor type, different methods are employed to change preset values for built-in interfaces. In the Series 300 family of SPUs the preset values can be changed by setting switches in a dual-in-line package (DIP switches). These switches may be found on the System Interface Board (or on the add-in card) on most Series 300 SPUs. On the models 345, 375, and 380, and on all Series 400 systems the select codes of internal (built-in) interfaces are *not* changed by physical switches, but by interacting with the Boot ROM Configuration program at system boot time. Details on this procedure is covered in the module entitled Boot Up (Series 300/400).



27

When installing add-in interfaces you should pay close attention to information provided in *HP-UX Installing Peripherals*.

## HW-13. SLIDE: Installation of Cards

14

### Electrostatic Discharge Precautions

- Read the card installation guide before you start
- Power off and disconnect all power cords
- Stand on a static-free mat
- Wear a static strap
- Connect all equipment together
- Keep uninstalled printed circuit boards in their protective antistatic bags
- Handle printed circuit boards by their edges

51436 HW-15. 5 

## Student Notes

Electrostatic charges can damage the integrated circuits on printed circuit boards. To prevent such damage from occurring, observe the following precautions during board unpacking and installation:

- Read the card installation guide before you start
- Power off and disconnect all power cords.
- Stand on a static-free mat.
- Wear a static strap. This strap is typically attached to your wrist at one end and the chassis of the computer on the other end. This ensures that any accumulated electrostatic charge will be discharged from your body to ground.
- Connect all equipment together, including the static-free mat, static strap, and peripheral units to ensure the the same ground reference point.
- Keep uninstalled printed circuit boards in their protective antistatic bags.

- Handle printed circuit boards by their edges, once you have removed them from their protective antistatic bags.

**Note** If you do not follow these precautions, you may cause serious damage to your system.

## HW-14. SLIDE: SCSI Cabling Guidelines

15

### SCSI Cabling Guidelines

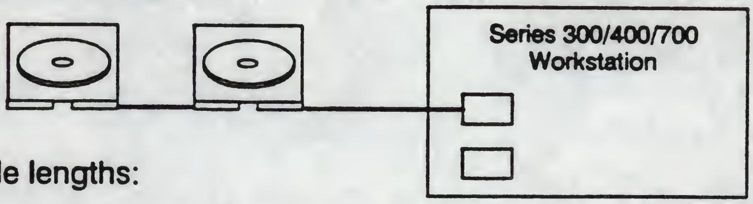


- Total cable lengths:
  - Single-ended SCSI II: 6 meters or shorter
  - Differential SCSI II: 25 meters or shorter
- Do not mix single-ended and differential SCSI II devices
- No cable segment less than 0.5 meters
- Maximum of seven SCSI devices per SCSI interface card
- Unique bus addresses between 0 and 6
- Terminator required on Host Adapter and last SCSI device in chain

51436 HW-16. 6 

## Student Notes

SCSI interface cards support multiple SCSI devices, what follows are some details and guidelines on how it's done.

The SCSI interface that comes standard on the Models 300s, 400s and 700s is a SCSI-II single-ended interface. (Single-ended SCSI II is also known as normal or standard SCSI II.) This interface supports data rates slower that the optional SCSI-II differential interface that plugs into a Model 700's EISA slot. SCSI-II single-ended supports data rates of 5 MB/sec synchronous and up to 5 MB/sec asynchronous. SCSI-II differential supports data rates of 10 MB/sec synchronous and up to 5 MB/sec asynchronous.

Single-ended and differential SCSI II devices cannot be mixed on the same interface.

Cable lengths of 6 meters and 25 meters are supported on single-ended and differential SCSI II respectively. No cable length segment can be less than 0.5 meters. Remember that devices internal to the chassis also have cable lengths that must be included in this maximum cable length calculation.



29

A maximum of seven SCSI devices are supported on a single SCSI interface card. Devices on these cards will each have a unique bus addresses between 0 and 6. SCSI bus address 7 is reserved for the SCSI interface and not to be used for a SCSI device bus address.

Terminators are required on each end of a SCSI chain namely on the host adapter (SCSI interface card) and last SCSI device in chain. The host adapter card in is terminated internally. For the last SCSI device in the chain, there are two types of SCSI terminators: high-density terminators and low-density terminators. These terminators are electrically the same, but mechanically different. In other words, if the high-density terminator won't fit on your device, you are safe if you use a low-density terminator and vice a versa.

## HW-15. SLIDE: SCSI Device Power Up Guidelines

16

### SCSI Device Power Up Guidelines

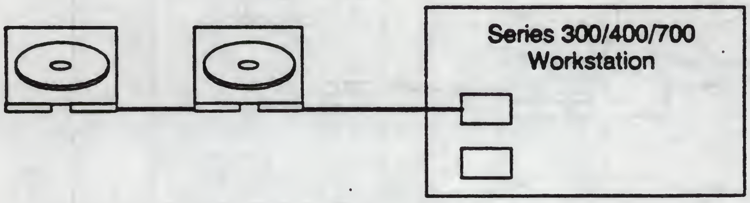


- *Do not* connect or disconnect any device while the system is running
- *Do not* turn power on or off to any device while the system is powered-up
- *Do* power on and complete self-tests on all peripherals before powering on SPU
- Change bus addresses with device powered off

hardware

51436 HW-17. 7 

## Student Notes

You will need to follow these SCSI device power up guidelines when changing the bus address on a SCSI device.

- *Do not* connect or disconnect any device while the system is running
- *Do not* turn power on or off to any device while the system is powered-up
- *Do* power on and complete self-tests on all peripherals before powering on SPU
- Change bus addresses with device powered off



30

## HW-16. SLIDE: HP-IB Device Guidelines

17

### HP-IB Device Guidelines

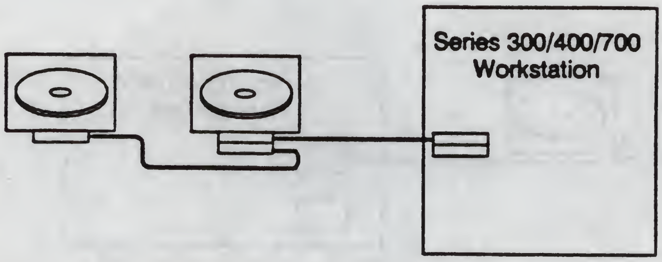


- *Do not* connect or disconnect any device while the system is running
- *Do not* turn power on or off to any device while the system is powered-up
- Change bus addresses with device powered off
- Do not mix high-speed devices with slower devices
- Maximum of eight devices

51436 HW-18. 8 

## Student Notes

HP-IB interface cards, like SCSI, support multiple HP-IB devices, what follows are some guidelines on how to power up these HP-IB devices. You will should follow these guidelines when changing the bus address on a HP-IB device.

- *Do not* connect or disconnect any device while the system is running
- *Do not* turn power on or off to any device while the system is powered-up
- Change bus addresses with device powered off. Since HP-IB addresses are only read at power up, a change in HP-IB bus address is ignored unless power is cycled on the device.
- Do not mix high-speed devices with slower devices
- Maximum of eight devices

## HW-17. SLIDE: Setting Device Addresses

18

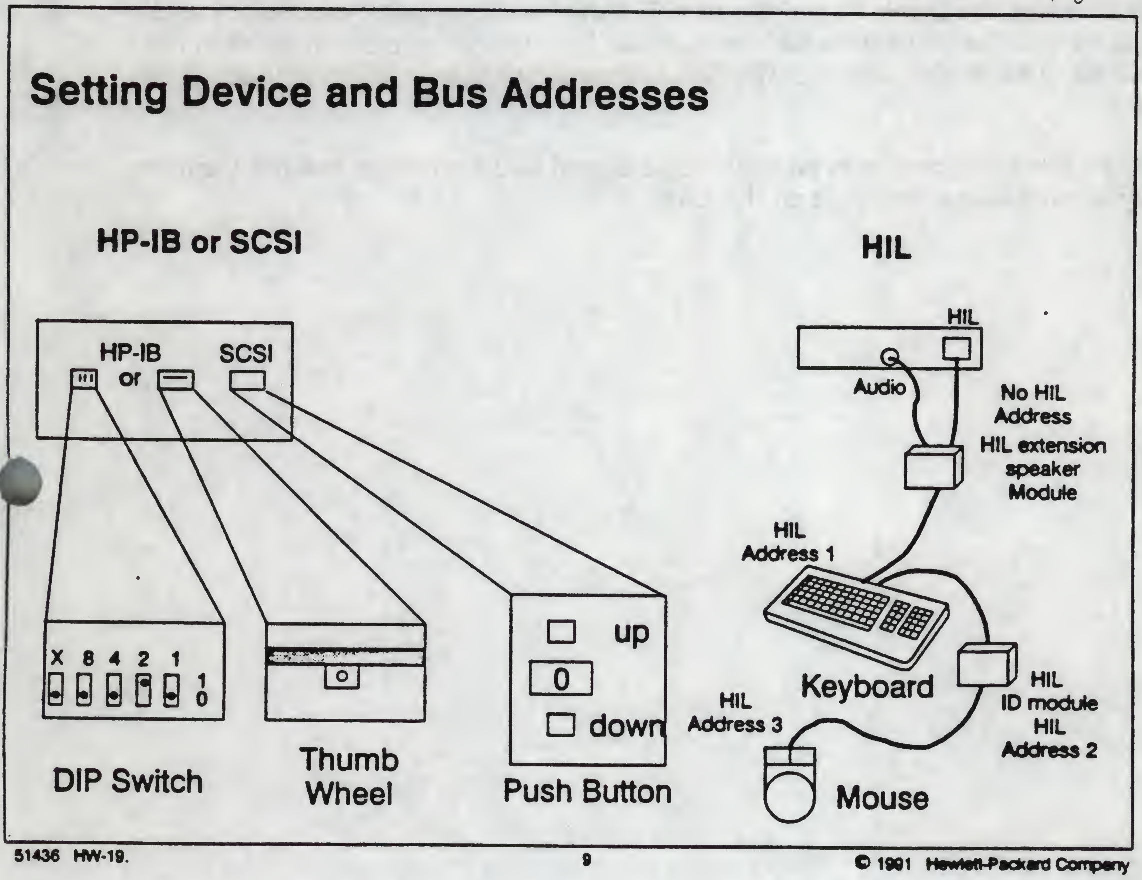


## Student Notes

If more than one device is physically connected to the same interface, the devices must be distinguished from one another by some kind of address. For SCSI and HP-IB devices, this address is known as the bus address. For HIL, this address is implied by the position of the device. For serial devices, the address is known as the port number.

On the back of each SCSI or HP-IB device, there is either a series of DIP switches or a thumb wheel switch with numbers on it. The settings of the DIP switches or the thumb wheel determine the bus address for the device. An alternate switch form is seen on some SCSI devices which involves a push button switch that increments or decrements the Bus Address.

There are suggested Bus Address settings for system disks that may be counter-intuitive. For SCSI disks, it is suggested that the system disk be at Bus Address six (6). For HP-IB devices, the system disk needs to be at HP-IB Bus Address zero (0). These choices are recommended due to the Boot ROM search sequence.



31

HIL devices do not have their HIL Bus address set manually with switches. The bus addresses are implicitly set by their position on the HIL bus. The device closest to the SPU is bus address one (1). This is usually the keyboard, but it need not be. An HIL ID module is usually placed between the keyboard and the HIL extension or between the keyboard and the next HIL device. Subsequent devices are at the next higher address, up to a limit of seven device addresses. Some HIL devices consume more than one HIL bus address (for example, 9 Knob dial module). The HIL Extension/Speaker modules, if used, do not consume an HIL address.

In the case of serial ports on the same card, port numbers are assigned to differentiate one from another. The addresses are usually labeled above each port on the card.

## HW-18. SLIDE: Connecting to the Local Area Network

19

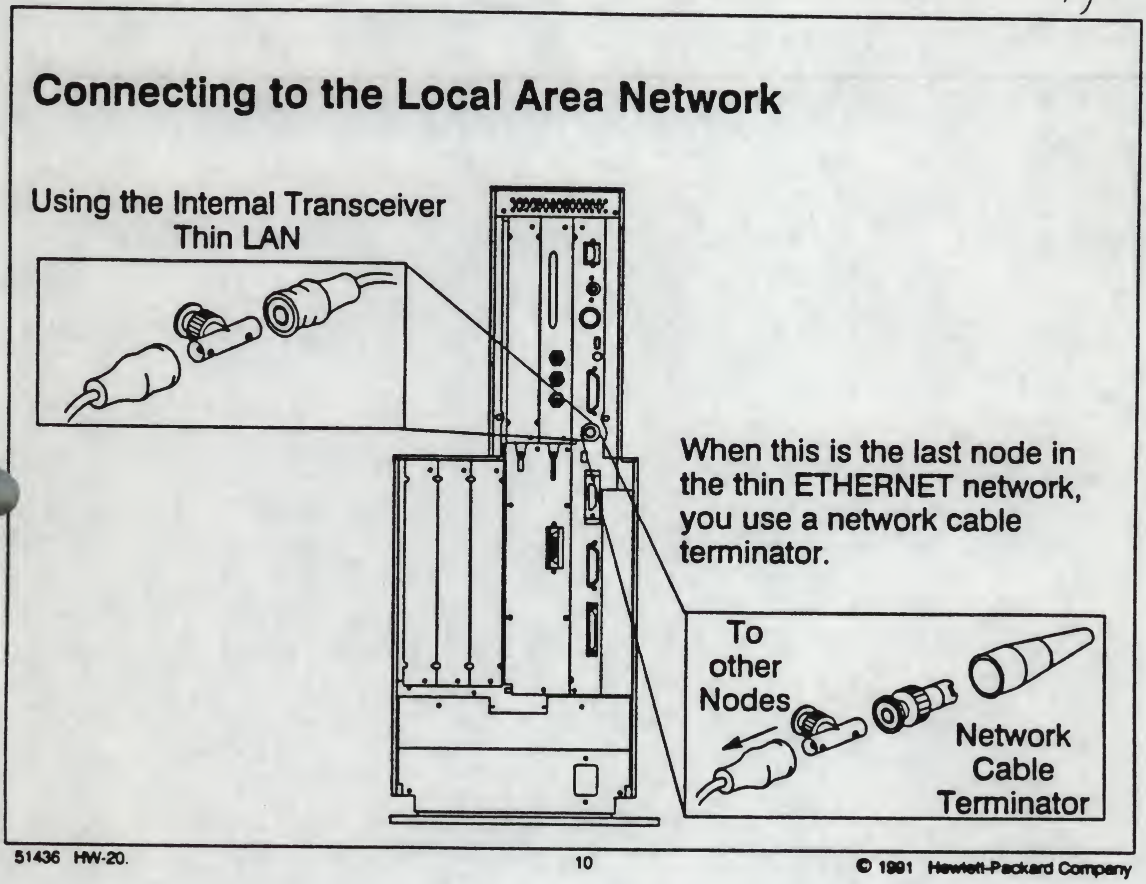


### Student Notes

A LAN connection is used for communication with other systems on a Local Area Network. The default (as shipped) is the BNC connection to a ThinLAN cable. In the case of the BNC connection, the Media Access Unit (MAU) or transceiver is integrated on to the LAN interface card. On the Series 700, the appropriate option needs to be order to convert this connection for use with a separate 15-pin AUI cable that supports a separate MAU to access ThickLAN backbone cable or EtherTwist cabling.

On the Series 300/400, you may choose to use the 15-pin AUI cable after you've already taken delivery of your hardware. To make this choice, move the internal jumper to select the AUI LAN port. See the installation manual for your SPU to perform this step. Unless you ordered a factory option to change the jumper, your system is shipped with a default setting for the use of the internal MAU or transceiver.

To connect to the internal transceiver (ThinLAN MAU), connect the BNC cables as indicated in the slide. The terminator may have a slightly different appearance than indicated in the slide. The slide depicts an insulated terminator.

**Note**

LAN cables require termination at the each extreme end of the cable.

## HW-19. SLIDE: Hardware for Backup and Installation

20

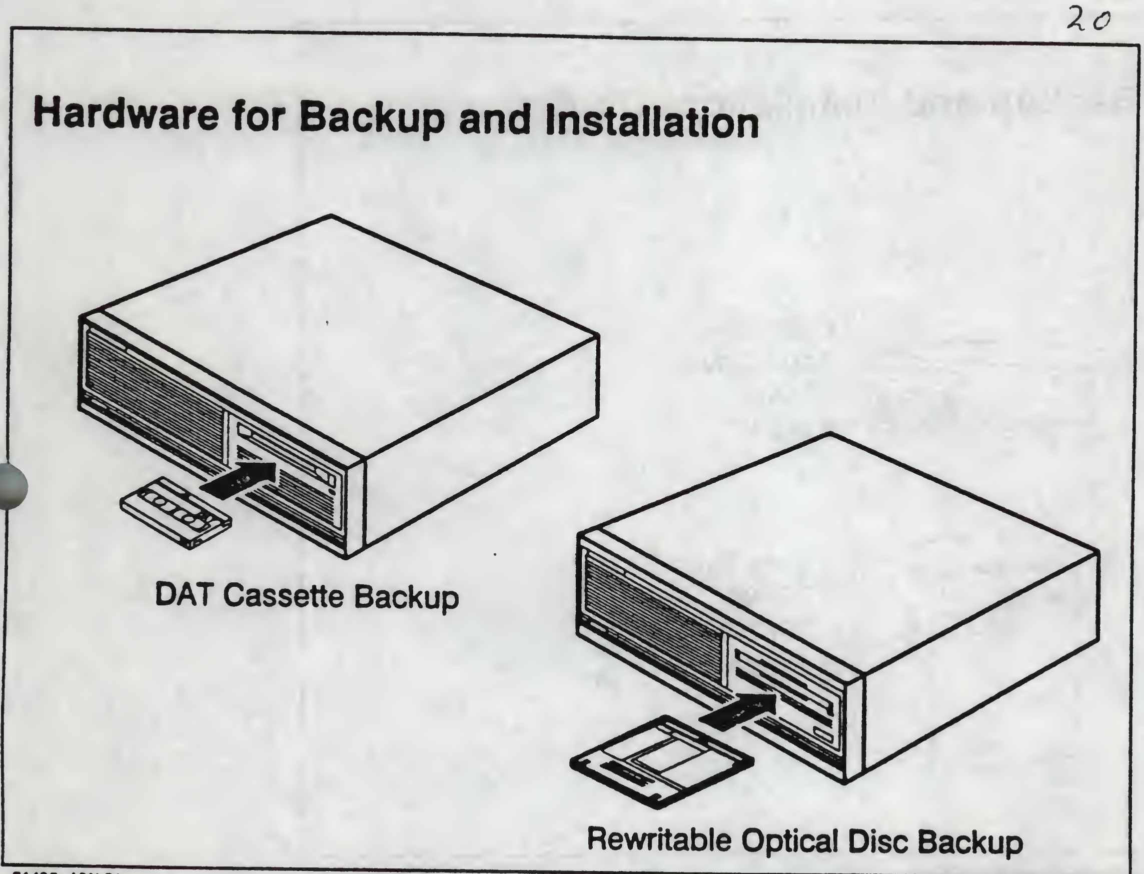


51436 HW-21. 11 

## Student Notes

Each system requires a means of backing up the HP-UX file system to prevent loss of data. Choices include:

- Digital Audio Tape (DAT) Digital Data Storage (DDS) Format tape drive The DDS (DAT) format tape drive is very high density (1.3 gigabytes per tape), low cost media and high storage transfer rate. It is supported for backup of the HP-UX file system on all HP-UX systems and an installation device on the Series 700 and some Series 800s. The DAT comes in a SCSI and an HP-IB version.
- Rewritable Optical Disk drive. The Rewritable Optical disk drive is sometimes called Magneto-Optical (MO) disk technology. MO disks are supported for backup of the HP-UX file system, and are a good choice due to high storage density and advantages as a removable file system. The rewritable optical disk stores 350 megabytes per side with two sides per cartridge for a capacity of 650 megabytes. This drive may either be SCSI or HP-IB.



33

## HW-19. SLIDE: Hardware Components for System Backup (Continued)

21

### Hardware for Backup and Installation (Cont'd)


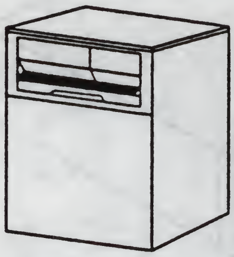


51436 HW-22. 12 

### Student Notes

Other types of backup devices that can be used to backup and recover files in HP-UX:

- Quarter inch cartridge tape drive.
  Cartridge tape, previously the standard HP distribution media for software, is still supported for update and installation at the 8.0 release. Cartridge tapes offer a variety of disk mechanisms, high reliability, low cost drive mechanisms, and a proven track record. Cartridge tapes do have lower storage density than the DDS or MO technology storage devices, but also have a convenient form factor and density superior to reel-to-reel tape drives. This may only be an external device with an HP-IB interface. (an autochanger version is not shown)
- 1/2 inch 9-track mag tape drive.
  Mag tape, although lower density than DDS, MO and Cartridge tapes is probably the most portable media between systems from different vendors due to its longevity in the computer industry. This portability advantage is likely to change as industry standards for DDS(DAT) format tapes and MO disks become widely accepted. This may only be an external device with an HP-IB interface.

## HW-20. SLIDE: CD-ROM

72

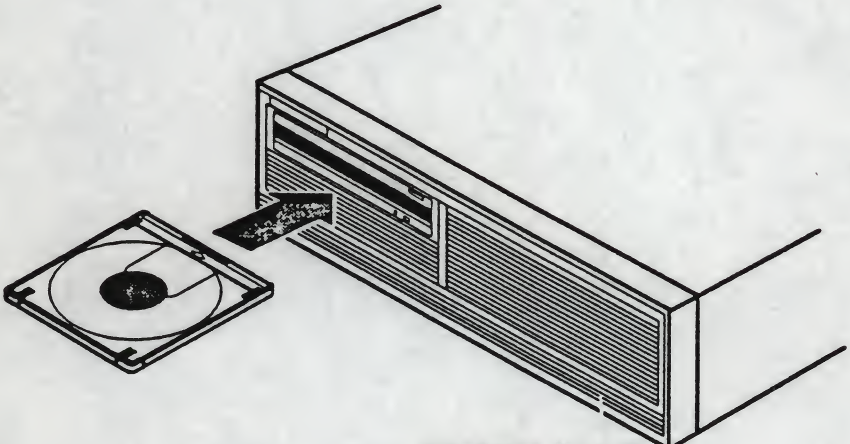


51436 HW-22. 13 

## Student Notes

CD-ROM is ideal for low-cost distribution of software, large volumes of data, documentation, and other information. Current data formats support 550 to 600 MB of data. Its random access and on-line data retrieval can go beyond the capability of magnetic tape, microfiche, or other distribution media.

CD-ROM is HP's recommended media for software installation. Since installation or updating is usually an infrequent activity, the CD-ROM drive may also be used for accessing system documentation in electronic form with the LaserROM/UX software. LaserROM/UX allows a friendly user interfaces to

- User and reference manuals
- Application Notes
- Software Status Bulletins
- HP Response Center Questions and Answers
- HP Quick Reference Product Catalogs
- Education Catalog
- Third-Party Software Catalog

**The HP LaserROM Information Service provides a new HP LaserROM disk every month.**

## HW-21. EXERCISES: Hardware (Series 700)

### Directions

Answer the following questions. Provide answers in both decimal and hexadecimal notation whenever appropriate.

1. What is the bus address of the system disk?

2. What is the SBM of the core I/O board?

3. What is the SBM of the EISA bus adapter?

4. What is the SBM of SGC (graphics)?

5. What is the function number of:

- SCSI
- Serial port 1
- Serial port 2
- LAN
- Parallel
- HIL

6. What is the bus address of the CD-ROM drive, if any?

## HW-22. EXERCISES: Hardware (Series 300/400)

### Directions

Answer the following questions. Provide answers in both decimal and hexadecimal notation whenever appropriate.

1. What is the default select code of the built-in high-speed HP-IB card in a HP 9000 Series 300 system to which the system disk drive is usually attached?

2. What is the default select code of the built-in SCSI interface in an HP 9000 Series 400 system to which the system disk drive is attached?

3. Assume a built-in SCSI interface in an HP 9000 Series 400 system connected to an internal disc drive. Is this interface different from the SCSI interface connector on the back panel? What is the select code for the rear panel interface?

4. What is the select code of the standard-speed built-in HP 9000 Series 300 HP-IB? (hint - a flexible disk drive or HP-IB printer may be attached to it.)

5. What is the select code of the optional standard-speed HP-IB available on some series 400s?



36

6. What is the select code of the built-in HP 9000 Series 300 serial port?

7. What is the select code of the parallel port on a Series 400? What is it on an HP 9000 model 375?

8. What is the bus address of the external disk drive your lab system? If an internal disk drive is used, typical on a Series 400, what is its bus address?

9. What is the bus address of the flexible disk drive, if one is used on your system? How many flexible disk drive units are present at that bus address?

10. What is the bus address of the HP-IB printer, if one is used on your system?

11. What is the bus address of the external HP-IB cartridge tape drive, if any?

12. What is the bus address of the external CD-ROM drive, if any?

37

# Module B3 — Boot Up (Series 300/400)

## Objectives

Upon completion of this module, you will be able to:

- List the different areas of the root disk.
- List the parts of the boot area.
- Boot up the Series 300/400 in attended and unattended modes.
- Configure the the Boot ROM menus, if applicable to your hardware.

# Module B3 — Boot Up (Series 300/400)

## Objectives

Upon completion of this module, you will be able to:

- List the different areas of the root disk.
- List the parts of the boot area.
- Boot up the Series 300/400 in attended and unattended modes.
- Configure the the Boot ROM menus, if applicable to your hardware.

# Module B3 — Boot Up (Series 300/400)

## B3-1. SLIDE: Root Disk Layout

23

### Root Disk Layout

Root disks can hold:

- Boot area (8KB)
- File system
- Swap

Kernel in

8K

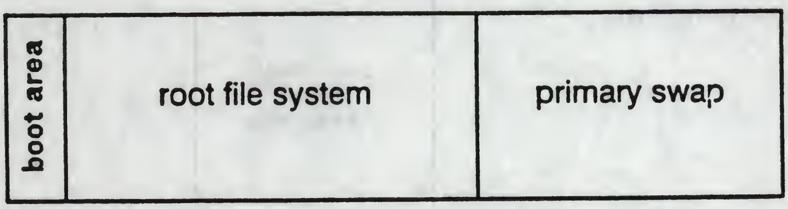


51436 B3-1. 1 

## Student Notes

On a Series 300/400 HFS implementation, a hard disk may be comprised of a boot area, a file system area, and a swap area. The **root disk** will hold all three of these areas. (The root disk or system disk is the disk that holds the root file system.) At the start of the root disk is a 8KB area reserved for the boot. This 8KB area is a file system in its own right, albeit not a a UNIX file system. Its format is Logical Interchange Format or LIF. LIF is a very simple file system with simple data structures that can easily fit in a small area such as 8KB. Its simplicity makes it a good choice at boot up in order to keep the code size small in the boot ROM. This LIF volume actually holds, the secondary loader, the boot up code that will load in the HP-UX kernel, `/hp-ux`. This area is followed by the file system area for the root file system. The primary swap area is located just after the root file system.

Non-root disks will typically contain a single swap area, a single file system, or a combination of a both a single swap area and a single file system area. On the Series 300/400 computer there is only one file system area per physical disk.



38

## B3-2. SLIDE: Boot Area

24

### Boot Area

- Volume header identifies disk format
- Volume directory contains entries for:
  - SYSHPUX - HUX Kernel.
  - SYSDEBUG [illegible]
  - SYSBCKUP Reserve
- Secondary loader:
  - Executes during system boot
  - Loads HP-UX system kernel

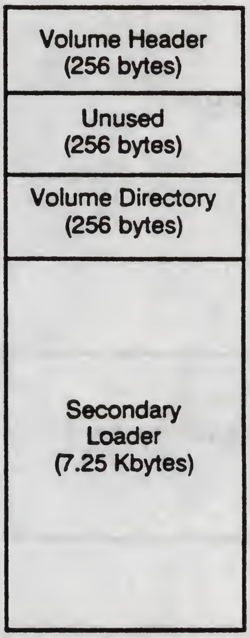


51436 B3-2. 2 

## Student Notes

The boot area is created on the root disk during the installation process. On a Series 300/400 system, the boot area contains a volume header, volume directory information, and a small secondary loader used when the system is loaded. The **volume header** in the boot area identifies the volume format.

The **volume directory information** contains three names: SYSHPUX, SYSDEBUG, and SYSBCKUP. SYSHPUX corresponds to the file /hp-ux, which is your kernel. SYSDEBUG corresponds to the file /SYSDEBUG. This file is used only when writing device drivers. SYSBCKUP corresponds to /SYSBCKUP, which is used as a backup kernel. All three are assumed by HP-UX to be object files.

The rest of the boot area contains the **secondary loader**. The boot ROM loads and passes control to the secondary loader which in turn loads and passes control to the file /hp-ux (or the backup kernel if you are using SYSBCKUP). /hp-ux then completes the task of booting HP-UX.

## B3-3. SLIDE: Boot ROM Startup Sequence

25

### Boot ROM Startup Sequence

The Boot ROM performs several tasks during system startup:

1. Locates and writes to a system console device.
2. Searches for and tests hardware internal to SPU.
3. Searches for an operating system.
4. Loads an operating system and starts it running.
5. Optionally configures internal interfaces in the SPU.

51436 B3-3. 3 

## Student Notes

The **Boot ROM** is a small machine-language program that resides in the ROM (Read-Only-Memory) located inside the SPU. When you boot (or simply, power-up) the system unit, the computer's processor executes (runs) the boot ROM program before any other. This program takes control of the system hardware and performs various tasks required to get the system running. The boot ROM performs these tasks:

- Locates and writes to a system console device.
- Searches for and tests hardware internal to SPU.
  During self-test, if some hardware fails a test or is improperly configured, the boot ROM displays an error message to that effect. For example, two interfaces set to the same select code will cause a boot ROM error.
- Searches an operating system.
- Loads an operating system and starts it running.



39

- Optionally configures internal interfaces in the SPU.
  The Boot ROM will enter an `EEPROM Configuration` menu if requested at boot time. This menu has three areas: You can Configure Boot Mode, Configure I/O Interfaces, or Configure Auto System Select mode The Boot ROM `EEPROM Configuration` menu is covered in more detail in a later slide.

You must have boot ROM Revision C or later (on an HP 9000 series 300) to boot from a SCSI disk drive. At Revision C, SCSI disk drives are treated the same as any other disk drive. You must have boot ROM Revision B or later (on a series 300) to boot across a LAN, as in a the HP-UX Clustered Environment.

The search order differs depending on the Boot ROM revision level. There is one search order for HP 9000 **Series 300** Boot ROMs with letter revisions (A, A2, B, C, C1, D), and a different order for Boot ROM revisions that are numbered revisions (1.1, 2.0, etc)

For HP 9000 **Series 300** with a Boot ROM revision D or earlier, the Boot ROM operating system search order is:

1. Any external mass storage interface (HP-IB or SCSI) at select codes 0-31 with bus address 0, unit number 0, and volume number 0. Search select codes in ascending order.
2. Shared Resource Manager (SRM) interface at select code 21.
3. Local Area Network (LAN) interface at select code 21.
4. Bubble Memory (HP-98259) Card at select code 30.
5. Other bootable mass storage interfaces not previously found. Search select codes in ascending order. Search bus addresses in ascending order. All bus addresses at each select code are searched first, before moving on to another interface at a higher select code.
6. SRM interfaces not previously found.
7. LAN interfaces not previously found.
8. Bubble Memory (HP-98259) interfaces not previously found.

For HP 9000 **Series 300 or Series 400** Boot ROM revision 1.1 or later, the Boot ROM searches for a bootable operating system differently. The Boot ROM operating system search order is:

1. SCSI mass storage interface at select codes 0-31 with bus addresses 7, 6, or 5, and unit number 0. Search select codes in ascending order. Search bus addresses 7, 6, or 5 in descending order.
2. HP-IB mass storage interface at select codes 0-31 with bus address of 0 and unit number of 0. Search select codes in ascending order.
3. Shared Resource Manager (SRM) interface at select code 14.
4. Local Area Network (LAN) interface at select code 21.
5. Bubble Memory (HP 98259) interface at select code 30.
6. Other SCSI bootable mass storage interfaces not previously found. Search select codes in ascending order. Search bus addresses in descending order. All bus addresses at each select code are searched first, before moving on to another interface at a higher select code.
7. Other HP-IB bootable mass storage interfaces not previously found. Search select codes in ascending order. Search bus addresses in ascending order. All bus addresses at each select code are searched first, before moving on to another interface at a higher select code.
8. SRM interfaces not previously found.

9. LAN interfaces not previously found.
10. Bubble Memory (HP-98259) interfaces not previously found.

If during a previous system boot, you used the Boot ROM EEPROM Configuration Menu, the Boot ROM's search order discussed above may not apply. If you made an `Auto System Selection` during EEPROM reconfiguration, then the specified disk address is assumed to be connected, and will be selected as the system disk without regard for the search sequence stated above. The `Boot ROM EEPROM Configuration` will be covered in detail later.



40

## B3-4. SLIDE: Series 300 Attended Mode Boot

### Series 300 Attended Mode Boot

```
Copyright 1989,                              :HP C1707 REMV, 700, 0,0
Hewlett-Packard Company.                        1H SYSHPUX
All Rights Reserved.                            1D SYSDEBUG
                                                1B SYSBCKUP
BOOTROM Rev. D                                  1T SYSTEST
Bit Mapped Video                             :HP 9145 TAPE, 703, 0,0
MC68030 Processor                               2H SYSHPUX
MC68882 Coprocessor                             2D SYSDEBUG
Configuration EEPROM                            2B SYSBCKUP
HP-HIL.Keyboard <-- press [Space] bar when you see this!
HP-IB
DMA-C0
RAM 16776942 Bytes
HP98644 (RS232) at 9
HP98265 (SCSI S 32) at 14
HP98643 (LAN) at 21, 080009ABCDEF
HP PARALLEL at 23
   .
   .
   .
SEARCHING FOR A SYSTEM (RETURN To Pause)
RESET To Power-Up
```

51436 B3-4. 4 

## Student Notes

The messages seen on the display at Boot time are shown on the slide for a model 375 system with Boot ROM revision D, your system may differ in the messages displayed and interfaces supported.

## Attended vs Unattended Booting

In **attended mode,** you select which operating system to boot. You choose among all the operating systems found by the search of mass storage by the Boot ROM. In **unattended mode,** you do not select which operating system to boot. The Boot ROM boots the first operating system it finds in the search of mass storage described earlier.

## Attended Mode Boot

You use attended mode when you want to choose a specific operating system (or if you do not want to boot from the first operating system found by the boot ROM). For example, when you are installing HP-UX, you need to boot from the install media, so you would use attended mode.

To enter attended mode, you must press any key after the keyboard has been initialized and before the boot ROM finds a default system and begins to load it. (The keyboard is initialized when the word `HP-HIL.Keyboard` or `Keyboard` appears on the system console.) Certain key combinations have special significance during the times described above, so we recommend pressing the Space bar to avoid unexpected behavior.

The attended mode boot process will now display available kernels to boot. These are shown on the right side of the graphic. Use the corresponding two letter abbreviation to select one of these kernels, for example `1H` for `SYSHPUX` on the CD-ROM.

## Unattended Mode Boot

In **unattended mode**, the boot ROM automatically boots the first operating system found. Unattended mode is the default method for loading an operating system. Use the unattended mode of booting if you have only one bootable operating system on-line, or if you know that the operating system you wish to boot will always be the first (or only) operating system the boot ROM will find.



41

## B3-5. SLIDE: Series 400 Attended Mode Boot

26

### Series 400 Attended Mode Boot

```
Copyright 1990,                                :HP C1707 REMV, 700,0
Hewlett-Packard Company.                          1H SYSHPUX
All Rights Reserved.                              1D SYSDEBUG
                                                  1B SYSBCKUP
BOOTROM Series 400 Rev. 1.1                    :HP 9145 TAPE,  703,0
MD12 REV 1.2 1990/08/07.14:27:08                  2H SYSHPUX
MC68030 Processor                                 2D SYSDEBUG
MC68882 Coprocessor                               2B SYSBCKUP
Configuration EEPROM
Utility Chip at 41
HP-HIL.Keyboard          <---- press [Space bar] when you see this!
HP-IB
DMA-C0
RAM 16776942 Bytes
HP 98644 (RS-232) at 9
HP PARALLEL at 12
HP 98265 (SCSI S 32) at 14
HP 98643 (LAN) at 21, THIN, 080009CBA89E
Bit Mapped Video at 133 (Console)
   .
   .
   .
SEARCHING FOR A SYSTEM (RETURN To Pause)
RESET To Power-Up
```

51436 B3-5. 5 

## Student Notes

The slide shows a typical series 400 computer system console display with Boot ROM messages following an "attended mode" boot. There may be slight variations in displayed messages from one boot ROM revision to another, for different series 400 SPU models, and for different installed interfaces.

## Attended vs Unattended Booting

In **attended mode**, you select which operating system to boot. You choose among all the operating systems found by the search of mass storage by the Boot ROM. In **unattended mode**, you do not select which operating system to boot. The Boot ROM boots the first operating system it finds in the search of mass storage described earlier.

## Attended Mode Boot

You use attended mode when you want to choose a specific operating system (or if you do not want to boot from the first operating system found by the boot ROM). For example, when you are installing HP-UX, you need to boot from the install media, so you would use attended mode.

To enter attended mode, you must press any key after the keyboard has been initialized and before the boot ROM finds a default system and begins to load it. (The keyboard is initialized when the word `HP-HIL.Keyboard` or `Keyboard` appears on the system console.) Certain key combinations have special significance during the times described above, so we recommend pressing the (Space) bar to avoid unexpected behavior.

The attended mode boot process will now display available kernels to boot. These are shown on the right side of the graphic. Use the corresponding two letter abbreviation to select one of these kernels, for example `1H` for `SYSHPUX` on the CD-ROM.

## Unattended Mode Boot

In **unattended mode**, the series 400 boot ROM will take one of several possible actions at boot time:

- Automatically boots the first operating system found
- Boot an operating system you specified by setting an "Auto System Selection."
- Display a menu for Configuration Control Selection



42

## B3-6. SLIDE: Series 400 Boot ROM Configuration Menus

27

### Series 400 Boot ROM Configuration Menus

```
Configuration Control

Keys   Mode  Class
---------------------------
 1  I/O Configuration
 2  Auto System Selection
 3  Boot Mode Selection

 A  Abort without changes
---------------------------
   Type [key] RETURN ?
```

51436 B3-6. 6 

## Student Notes

Some Series 300 and all Series 400 Boot ROMs allow you to change the built-in I/O interface configuration by interacting with this interacting with this on-screen menu. Some Series 300 SPUs support a System Interface Board described in the hardware overview module. We described DIP switches on that board when we talked about setting or changing the Select Code and configuration for built-in interfaces. Some Series 300 and all Series 400 systems do not have actual DIP switches, as such we use this boot ROM capability to change the default built-in I/O Interface Configuration. Normally you will not have to change the select code, interrupt level, or other configuration settings for built-in interfaces.

The slide shows the Boot ROM I/O Interface Configuration menu for an HP 9000 model 375. A "Configuration EEPROM" is supported on only a few series 300 models, but is supported on all series 400 models. *Only* if the Boot ROM messages at boot time include "Configuration EEPROM", is this feature supported on your system processor.

The boot ROM on all Series 400 models offers three configurable areas as we saw in the last slide. These are Boot Mode Configuration, Interface Configuration, and Auto System Select Mode.

# Boot ROM I/O Configuration Menu

```
Copyright 1989,
Hewlett-Packard Company.
All Rights Reserved.

BOOTROM Rev. D
Bit Mapped Video
MC68030 Processor
MC68882 Coprocessor
Configuration EEPROM
HP-HIL.Keyboard
HP-IB
DMA-C0
RAM 16776942 Bytes
power)
HP98644 (RS232) at 9
HP98265 (SCSI S 32) at 14
HP98643 (LAN) at 21, 080009ABCDEF
HP PARALLEL at 23
    .
    .
    .
System Search Mode
RESET To Power-Up, Space to clear input
```

```
     Configurable Interfaces
Keys Interface      Select Code
-------------------------------
  1  HP-IB
  2  RS-232                   9
  3  SCSI                    14
  4  HS HP-IB                15
  5  LAN                     21
  6  HP PARALLEL             23

  N  store New values
  D  store Default values
     (then cycle SPU

  A  Abort without changes
 ------------------------------
 Type [key] RETURN ?
```




4.3

# Module B3 — Boot Up (Series 300/400)

The boot ROM on models 345 and 375 offers only one configurable area: `I/O Interface Configuration` which is automatically entered, when requested by the key sequence "CTRL-C" (Return) at boot time. This key sequence is used instead of the space bar to enter attended mode.

If you select I/O Interface Configuration, you will see a display at attended mode boot time as indicated in the slide. Of course the left hand side of the display depends upon the SPU type (300 or 400), and the list of configurable interfaces depends upon which interfaces are installed in your SPU. As an example, on a model 400s without an optional HP-IB, that entry will not appear in the list.

As you select various interfaces to configure, each selection displays a separate menu. These menus are specific to each kind of internal interface and have choices that apply only to the configuration of that interface. Once made, the choices are stored in EEPROM until explicitly changed in a later reconfiguration. The EEPROM is a read only memory that may be electrically erased and reprogrammed. As an Electrically Eraseable Read Only Memory, its contents are "remembered" even if power is removed from the System Unit, just as the contents of the Boot ROM are retained.

Some of the choices from the interface configuration menus illustrate the range of configurable interface characteristics. These are listed below.

Select Code - set value for most interfaces

Interrupt Level - set value for most interfaces

Remote/Local - make a serial terminal the console

Fast/Normal - set data transfer or handshake speed

Modem - enable modem handshaking for serial ports

Parity - parity checking of data for SCSI transfers

Bus Address - set value for SCSI Controller Bus Address

HP-IB system controller - enable/disable only

DMA mode - set the width of data transfers

---

**Note** Not all of these settings apply to every interface. You should read the reference manuals and system documentation for specific settings that apply to a given interface.

---

The standard speed HP-IB interface select code cannot be reconfigured, but other switch settings affecting the HP-IB may be changed.

Some references for reconfiguring your internal interfaces are *BootROM Configuration Mode Users Manual, HP 9000 Products*, and *HP-UX Installing Peripherals, HP 9000 series 300/400.*



64

## B3-7. LAB: Boot Up (Series 300/400)

### Directions

Complete the following questions.

1. Boot the kernel.

2. Boot the backup kernel.

3. Do an attended mode boot on your system. Determine if your lab system supports Boot ROM Configuration.

4. If your lab system supports Boot ROM Configuration, enter the I/O Configuration menu. Record the assigned Select Codes for the SPU interfaces. Do not reconfigure the I/O, just record the current mapping of select codes to interfaces.

5. If your lab system is an HP Apollo 9000 series 400, Enter the Boot ROM Configuration menu. Determine if an Auto Selected System has been assigned to this hardware. If HP-UX has not been selected, assign that value permanently in the EEPROM. If it has already been selected do NOT change it.

45

45

# Module B7 — Boot Up (Series 700)

## Objectives

Upon completion of this module, you will be able to:

- List the different areas of the root disk.
- List the parts of the boot area.
- Boot the Series 700 in attended mode.
- Boot the Series 700 in unattended mode and interact with:
  - Boot-up main menu
  - Boot Console User Interface

# Module B7 — Boot Up (Series 700)

## B7-1. SLIDE: Root Disk Layout

### Root Disk Layout

Root disks hold:

- LIF directory (8KB)
- File system
- Swap
- Boot area (2MB)

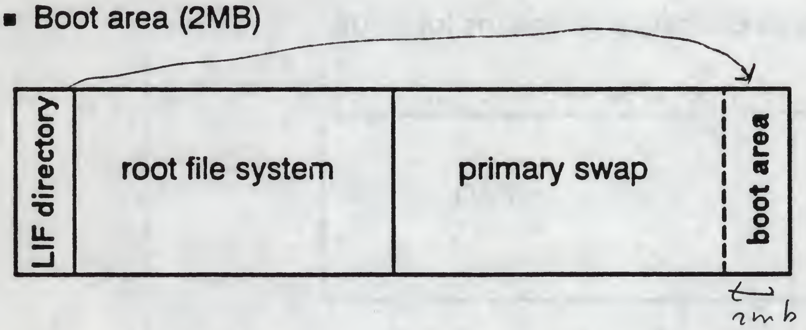


51436 B7-1. 1 

## Student Notes

On a Series 700 HFS implementation, a root disk may be comprised of a LIF directory, a file system area, a swap area, and a boot area. The root disk will hold all four of these areas. (The root disk or system disk is the disk that holds the root file system.) At the start of the root disk is a 8KB area reserved for the boot.

This 8KB area is a file system in its own right, albeit not a a UNIX file system. Its format is Logical Interchange Format or LIF. LIF is a very simple file system with simple data structures that can easily fit in a small area such as 8KB. Its simplicity makes it a good choice for boot up code in order to keep the code size small in the boot ROM. On a Series 700, this LIF area is actually a directory to direct the boot up process to the last 2MB on the root disk that actually holds, the bootstrap utility, the boot up code that will load in the HP-UX kernel. This 2MB area is actually taken from the last 2MB of swap on the root disk. Swap will occupy the end of this disk save the last 2MB. The file system area is located between the 8KB LIF directory and swap for the root disk.

## B7-2. SLIDE: Boot Area

### Boot Area

The boot area, residing on then last 2MB of the root disk, contains:

| | |
|---|---|
| `ISL` | Operating system independent Initial Program Loader (IPL) |
| `HPUX` | HP-UX specific initial system lcader utility for bootstrap and first-time installation |
| `AUTO` | Autoexec file containing operations and options for `HPUX` |

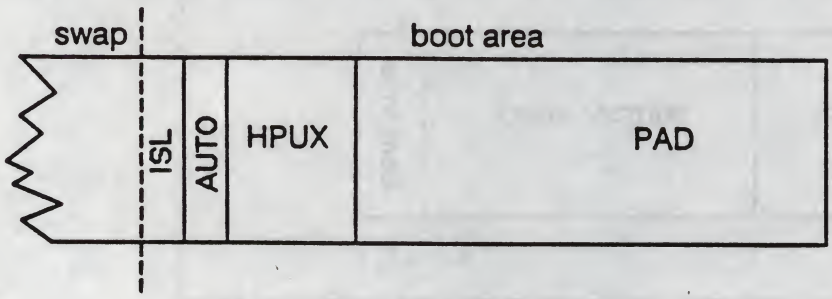


51436 B7-2. 2 

## Student Notes

At boot up, the Series 700's Processor Dependent Code (PDC) provides the functionality for booting the machine. PDC initializes the hardware and selects a console and boot device, and then loads the the initial program loader, ISL. ISL, for hardware that supports HP-UX, will invoke HPUX. HPUX will use the contents of AUTO to figure out what to boot.

For the Series 700, the boot area is automatically copied by the install process from the tape to last 2MB of the root disk.

ISL implements the operating system independent portion of the bootstrap process. It is loaded and executed after self-test and initialization have completed successfully.

AUTO is the autoexec file that may contain operations and options for HPUX. Although in actual practice on the Series 700, this file will usually contain nothing but the hpux command by itself.

# Module B7 — Boot Up (Series 700)

HPUX is the HP-UX specific initial system loader utility for bootstrap and first-time installation. HPUX is invoked from the ISL> prompt and is capable of executing these commands:

| | |
|---|---|
| `boot` | Loads an object file from an HP-UX file system, Logical Interchange Format (LIF) file system, or raw device, then transfers control to the loaded image. |
| `ls` | Lists the contents of HP-UX directories in a format similar to `ls(1)` |
| `restore` | The `restore` operation is provided as a recovery mechanism in the event that a disk becomes totally corrupted. We'll talk about how to restore a corrupted disk in the upcoming material. |
| `-v` | Lists the revision of the hpux loader. |

## B7-3. SLIDE: Boot Console User Interface: Main Menu

### Boot Console User Interface: Main Menu

```
32 MB of memory configured and tested.

Selecting a system to boot.
To stop selection process, press and hold the ESCAPE key.

Selection process stopped.

Searching for Potential Boot Devices.
To terminate search, press and hold the ESCAPE key.

Device Selection      Device Path              Device Type
-------------------------------------------------------------------------

P0                    scsi.6.0                 QUANTUM PD210S
P1                    scsi.5.0                 QUANTUM PD210S

    b)     Boot from specified device
    s)     Search for bootable devices
--> a)     Enter Boot Administration mode
    x)     Exit and continue boot sequence
    ?)     Help

Select from menu:  b p0 ipl
```

51436 B7-3. 3 

## Student Notes

Each time the machine is powered on, the user is given the opportunity to interact with the machine at a variety of places. For unattended mode, the user need never interact with the machine.

If the user wishes to interrupt the boot sequence to redirect the boot sequence or perform one of the boot administration functions provided by the Boot Console User Interface, the user should hit the ESCAPE key until the message,

`Terminating selection process.`

is displayed.

The automatic boot sequence has now been halted and the user is in fully "attended" or interactive mode.

At this point, the Boot Console User Interface is invoked and the first message displayed:

# Module B7 — Boot Up (Series 700)

```
Searching for potential Boot devices.
```

At this point, the user may do nothing and wait for the list of supported boot devices to be displayed. This may take some number of seconds as the SCSI, LAN, and EISA interfaces are searched.

```
Device Selection      Device Path              Device Type
-------------------------------------------------------------------------

P0                    scsi.6.0                 QUANTUM PD210S
P1                    scsi.5.0                 QUANTUM PD210S

b)    Boot from specified device
s)    Search for bootable devices
a)    Enter Boot Administration mode
x)    Exit and continue boot sequence
?)    Help

Select from menu:
```

The user may type the arguments to the b(oot) command at this menu. Examples follow illustrating booting up from the disk at bus address 6. The second example is an interactive boot.

```
Select from menu:  b p0

Select from menu:  b p0 isl
```



4P

## B7-4. SLIDE: Boot Console User Interface: BOOT_ADMIN

### Boot Console User Interface: BOOT_ADMIN

```
Select from menu:  a

BOOT_ADMIN> info

------------------------- Hardware Configuration --------------------

Machine model: 9000/720

Processor Frequency          = 50000000 Hz
I/O Subsystem Frequency      = 25000000 Hz

SCSI Jumper Frequency setting = 50000000 Hz

LAN Jumper Status: Internal ThinLAN Port selected

Processor Revision                    = 2
System Controller Revision            = 1
Floating Point Coprocessor Revision = 3


Hardware Version           8192     (0x00002000)
Software Version           1153     (0x00000481)
```

51436 B7-4 4 

## Student Notes

The Boot Console User Interface main menu provides access to the BOOT_ADMIN prompt. This interface will allow you alter default behaviors exhibited at boot up and also obtain useful information about the the hardware as configured.

From the Boot Console User Interface main menu type:

```
Select from menu:  help a
```

This will display:

```
The following commands are available:

AUTO                Display state of Autoboot/Autosearch flags
AUTOSEARCH          Set state of Autosearch flag
AUTOBOOT            Set state of Autoboot flag
```

```
BOOT                 Boot from Primary/Alternate path or Specified Device
DATE                 Read/Set the Real-Time Clock
EXIT                 Return to previous menu
FASTSIZE             Display/Set FASTSIZE memory parameter
HELP <item>          Display Help information for <item>
INFO                 Display boot/revision information
LAN_ADDR             Display LAN Station Address
OS                   Display/Select Operating System
PATH                 Display/Modify Path Information
PIM_INFO             Display Processor Internal Memory Information
RESET                Reset the System
SEARCH               Search for boot device
SECURE               Display/set secure boot mode
SHOW                 Display the results of the previous search
```

These boot admin commands are easy to use and display useful information, for example:

```
Select from menu:  a

BOOT_ADMIN> info

----------------------------- Hardware Configuration ------------------------

Machine model: 9000/720

Processor Frequency          = 50000000 Hz
I/O Subsystem Frequency      = 25000000 Hz

SCSI Jumper Frequency setting = 50000000 Hz

LAN Jumper Status: Internal ThinLAN Port selected

Processor Revision                   = 2
System Controller Revision           = 1
Floating Point Coprocessor Revision  = 3


Hardware Version             8192     (0x00002000)
Software Version             1153     (0x00000481)
```



49

## B7-5. SLIDE: hpux boot

hpux boot

hpux [ -*istring* ] boot [ *devicefile* ] [ *arguments* ]

Examples:

Booting up the kernel:

```
ISL> hpux boot disk(scsi.6;0)/hp-ux
```

Booting up the back-up kernel:

```
ISL> hpux boot disk(scsi.6;0)/SYSBCKUP
```

Booting up the kernel in single user mode:

```
ISL> hpux -is boot disk(scsi.6;0)/hp-ux
```

51436 B7-5. 5 

### Student Notes

ISL> is the interactive boot prompt. To access this interface ask for isl from the main menu. For example:

```
Select from menu:  b p0 isl
```

From isl, the hpux boot, hpux restore, and hpux ls operations accept *devicefile* specifications, which have the following format:

*manager(path;n)filename*

| | |
|---|---|
| *manager* | Valid *managers* are disk, tape, and lan managers. disk manages all disks. Lan manages remote boot through the LAN connection. Tape manages the DAT tape drive. |
| *path* | For the core I/O board, valid managers are scsi and lan. |
| *minor* | Currently all minor number are 0. They are reserved for future enhancements. |

*filename* Specified file names can be standard HP-UX path names.

The **hpux** boot command also supports an HP-UX argument parameter. `-i`*string* accepts a string that specifies the initial run-level for `init`. This will override the default run-level which is typically set to multi-user or run-level 2 in the `/etc/inittab` file.



50

## B7-6. SLIDE: hpux restore

hpux restore

The restore operation is provided as a recovery mechanism in the event that a disk becomes totally corrupted.

To create a restore dd tape:

```
# dd if=/usr/lib/uxbootlf.700 of=/dev/rmt/0mn bs=2k
# dd if=/dev/rdsk/6s0 of=/dev/rmt/0m bs=64k
```

To restore:

```
Select from menu:  b p2 ipl

ISL> hpux restore disk(scsi.6;0)
```

51436 B7-6. 6 

### Student Notes

The **restore** operation is provided as a recovery mechanism in the event that a disk becomes totally corrupted. To create the dd to restore from can be created from the usual root login when the system is in single user mode.

```
# dd if=/usr/lib/uxbootlf.700 of=/dev/rmt/0mn bs=2k
# dd if=/dev/rdsk/6s0 of=/dev/rmt/0m bs=64k
```

The first dd command puts a boot area on the DAT. The second copies the root file system to the DAT.

The following illustrates how to boot from a "bootable" DAT. A bootable DAT is one that has a valid LIF directory and boot area. The above example shows how to create a bootable DAT. Another example of a bootable DAT is the INSTALL tape created by HP.

```
Device Selection     Device Path                Device Type
---------------------------------------------------------------------
```

```
P0                  scsi.6.0                 QUANTUM PD210S
P1                  scsi.5.0                 QUANTUM PD210S
P2                  scsi.4.0                 HP      HP35450A    -A
P3                  scsi.1.0                 HP      2213A

b)    Boot from specified device
s)    Search for bootable devices
a)    Enter Boot Administration mode
x)    Exit and continue boot sequence
?)    Help

Select from menu:  b p2 ipl
```

The next command is specific to restoring a dd of the system disk. This command will restore a dd backup to a corrupted scsi disk supported off of the core I/O board at bus address 6 from a DAT at bus address 4.

```
ISL> hpux restore disk(scsi.6;0)
```

## B7-7. LAB: Boot Up (Series 700)

1. Boot the kernel by interacting with the Boot Console User Interface.

2. Boot the kernel in single-user mode.

3. Boot the kernel from the `BOOT_ADMIN>` interface.

4. Find the LAN station address.

5. What is the default boot path.

6. What is the difference between `BOOT_ADMIN> search` and `BOOT_ADMIN> show`? When would you use each?

7. Create and boot a dd restore tape.



52

7. Create and boot a dd restore tape.



52

# Module IH — Installing HP-UX (Series 300/400/700)

## Objectives

Upon completion of this module, the student will be able to do the following:

- Differentiate between an installation and an update.
- Perform the steps to install an HP-UX system on either an HP 9000 series 300, series 400, or series 700 system.

# Module IH — Installing HP-UX (Series 300/400/700)

## IH-1. SLIDE: HP-UX Installation

28

### HP-UX Installation

- Destroys any data that previously existed on the root disk.
- Constructs an HP-UX file system on the root disk.
- Copies subsystem files from the installation media to the root disk.
- Builds the initial HP-UX kernel on the root disk.
- Creates a login for the system administrator.

HP-UX installation devices include:

| | |
|---|---|
| Series 300/400: | Cartridge tape and CD-ROM |
| Series 700: | DDS format DAT and CD-ROM |

51436 IH-1. 1 

### Student Notes

Once you have properly installed your hardware, you are ready to install the software. Whenever a system installation needs to be done, the HP-UX system administrator is the only individual who should be permitted to do it. A system installation is required when a new system is ordered with out pre-loaded software or "instant ignition". A system installation may need to be done if there is a catastrophic failure with the system. Installation assumes that there is no data of value existing on the disk. Installation is not required when a new release of HP-UX is offered, usually an update is performed.

The installation process destroys any data that previously existed on the root disk. Next, an HP-UX file system is constructed on the root disk.

Subsystem files from the installation media will then be copied to this newly constructed file system. An HP-UX kernel will be created. This kernel may need to be customized to accomadate the applications that will run on this system. Kernel customization is covered in the module entitled Module RK. Lastly, a login is created for the system administrator.



53

# Module IH — Installing HP-UX (Series 300/400/700)

It is very important to understand that the process of installing HP-UX is different from updating HP-UX. The update process starts with a working HP-UX system and then modifies the existing system and files. By contrast, the install process overwrites any contents on the disk. During the course of the installation, you will interact with various programs. One of these programs is actually the update program.

The installation media is a cartridge tape or CD-ROM disk for a Series 300/400 system. For a Series 700 system, the installation media is a CD-ROM disk or DDS Format DAT.

## IH-2. SLIDE: Before You Start an Installation

29

### Before You Start an Installation

- Verify that system hardware has been correctly installed.
- Read the *README BEFORE INSTALLING HP-UX* notice.
- Determine:
  - The hardware address of the installation device and the root disk.
  - The amount of swap space needed.
  - Any required codewords and hardware IDs.
  - Which filesets and partitions you wish to install.
  - Customizing information (time zone, time, host name, IP address)

51436 IH-2. 2



### Student Notes

Before You Start an install:

- Verify that system hardware has been correctly installed.
- Read the *README BEFORE INSTALLING HP-UX* notice.
  If there was an installation manual supplied with your installation media, read it as well. Other excellent manuals for installation are *Installing and Updating HP-UX HP 9000 series 300/400* or *Installing and Updating HP-UX HP 9000 series 700.*
- Determine the hardware address of the installation device and the root disk.
  Determining these addresses is not absolutely necessary as the boot and installation processes will search for possible sources and destinations. However, knowing these addresses will help you confirm the correct device addresses as the installation program runs.
- Determine the amount of swap space needed.
  See the swap modules on techniques and considerations for determining this space.



54

- **Determine any required codewords and hardware IDs.**
  Starting at the HP-UX 8.0 release, with the introduction of CD-ROM software distribution, codewords have been added as security measure to the installation process. The codewords are based on specific software products and are tied to specific computer system hardware IDs. Codewords will be supplied on a software product certificate that you received with your distribution media. The hardware IDs are derived from the specific computer system that will run the software you are installing. At this time only installations using CD-ROM media require the use of codewords. Codewords and hardware IDs are covered in detail in Module UH.

- **Determine which filesets and partitions you wish to install.**
  If you are not planning to install everything from the software bundle that you ordered, determine which filesets and partitions you wish to install. A partition is a group of filesets, and a fileset is a group of files.

- **Determine customizing information (time zone, time, hostname, IP address)**
  If you plan to use networking, you will need to know a unique hostname for your system and a unique Internet Protocol (IP) address for your LAN interface. This information may not be available at the time you install the system, so there are defaults provided for you. Hopefully, time zone and time should be obvious and easily obtainable at install time.

## IH-3. SLIDE: Steps to Install the HP-UX Operating System

30

### Steps to Install the HP-UX Operating System

1. Turn on power to the disk drive, installation device, monitor, and expander (if applicable).
2. Insert the media labeled *INSTALL* into the installation device. If using tape, wait until the "busy" light goes off.
3. Turn on SPU.
4. Boot in attended mode

51436 IH-3. 3 

## Student Notes

### Using CD-ROM Install Media

Before inserting an *INSTALL* CD-ROM disc into the CD-ROM drive, make sure that the drive is powered ON and the front panel lights are out.

You should not have to force the disc carrier into the drive, but you will need a firm, steady push. Hold the disc carrier at the end labeled "Compact Disc", and push the carrier evenly and firmly into the drive until it bottoms-out. Then you should push it a bit further, release it and the carrier will pop back out partially, leaving the CD-ROM disc inside the drive. Pull the now empty carrier fully out, and set it aside, you will need it later to extract the CD-ROM disc in the drive.

## Using Cartridge Tape Install Media

Before inserting an "install" tape into the cartridge tape drive, check to make sure that it is write protected or "safe". The arrow in the top left corner should be turned to point to "SAFE."

## Using DDS Format DAT Install Media

Before inserting an "install" tape into the tape drive, check to make sure that it is write protected or "safe". The white tab should be slide to the open position.

## IH-4. SLIDE: Installation Attended Mode Boot Up (Series 300/400)

31

### Installation Attended Mode Boot Up

### (Series 300/400)

Turn on your SPU.

Configure "Boot Mode Selection" on Series 400, if necessary

**Press the space bar immediately** and hold it down

Release space bar when the `Keyboard` message appears.

Choose installation media as boot device.

Interact with menus presented by installation process.

51436 IH-4. 4 

## Student Notes

Once you turn the SPU power ON, your system will begin to boot up. The boot ROM will search for a system console, perform hardware tests, and then begin a search for an operating system. On a Series 400 that has never been powered up, the "Boot Mode Selection" menu needs to be configured for HP-UX. Details on the "Boot Mode Selection" are covered in the Boot Up (Series 300/400) module. A quick review on attended mode boot up follows.

To perform an attended mode boot, **Press the space bar immediately** and hold it down until you see: `HP-HIL.Keyboard` or `Keyboard` on the screen. A menu will now be displayed to allow you to choose an operating system. Enter the **two letter character code** that represents `SYSHPUX` on the Installation device that matches your choice of media. Notice that the program requires that the letters are entered as uppercase on a terminal and lowercase on an ITE (which does case folding for us). A (Return) is optional on the terminal and the on the ITE.

You should now read messages presented, follow the directions given interactively, and respond to the menus in the installation utility.



56

**Note**

If you see messages starting in the lower left of the screen and moving upwards, you have improperly selected the operating system in `Boot ROM Configuration`. You must start over by cycling the SPU power off and then on again. You must re-enter the Boot ROM's Configuration mode to correct this problem. See the owners manual for help on this.

The rest of the installation consists of answering questions and making choices on menus. The installation utility will take you through the process of installing your system. We will take a look at some of these menus on the following slides.



58

## IH-7. SLIDE: System Type Menu (Series 700)

### System Type Menu

```
HP-UX INSTALLATION UTILITY — SYSTEM TYPE MENU

This screen lets you choose whether or not you want a long-filename
system. You may convert to a long-filename system at any future time,
but once you have a long-filename system you can't go back to a
short-filename system.

The current choice is shown. Enter "y" (yes) or "n" (no)
to change a choice. Press "Done" when finished choosing.

          Long file names:    Y

[ ] [ ] [ ] [Done] [ ] [ ] [ ] [ ]
```

51436 IH-7. 7 

## Student Notes

HP-UX file systems can be configured to support either long or short file names. The standard file name length is 14 characters. Optionally, this limit can be increased to 255 characters. When configured with the standard limit, HP-UX file systems are compatible with earlier system releases that are not configured to accept long file names.

Generally, you install the system with long file names (the 255 character limit) to gain flexibility in naming files. Also, files created on other systems that allow long file names can be moved to your system without being renamed.

Avoid long file names if:

- You plan to use applications that read directory file information and do not use portable directory routines like those described in `directory(3C)`. If these applications assume that directories are in an array of fixed-size entries, they will not work with long file names. To correct this, rewrite the

## IH-5. SLIDE: Installation Attended Mode Boot Up (Series 700)

### Installation Attended Mode Boot Up

### (Series 700)

Turn on your SPU.

**Press the escape key immediately** and hold it down

Release escape key when the `Terminating selection process` message appears

Choose installation media as boot device.

Interact with menus presented by installation process.

51436 IH-5. 5 

## Student Notes

Once you turn the SPU power ON, your system will begin to boot up. The boot ROM will search for a system console, perform hardware tests, and then begin a search for an operating system.

To perform an attended mode boot, **Press the escape key immediately** and hold it down until you see: `Terminating selection process` appears on the screen. A menu will now be displayed to allow you to choose an operating system. Enter the **two letter character code** that represents the installation device that matches your choice of media.

You should now read messages presented, follow the directions given interactively, and respond to the menus in the installation utility.



57

## IH-6. SLIDE: Welcome Menu

32

### Welcome Menu

```
Welcome to HP-UX install.  There are basically 4 steps to installing HP-UX,
which this and other utilities will lead you through.

     Step 1) Select basic system characteristics, then specify
             the "destination disk" which is to hold HP-UX

     Step 2) Optionally  modify the file system parameters pre-set
             for your chosen destination disk.

     Step 3) Optionally change the amount of swap space on the
             destination disk.

     Step 4) Choose the filesets (functional groups of files) which
             you want loaded onto the destination disk.

Special function keys will be set for the menus used by each
step.  The arrow keys on the keyboard allow ycu to move the highlight
through menu items.  "Select Item" and the Return key act on menu
items; "Return" is also used to enter values typed at the
keyboard.

Press "Return" when you're ready to proceed to Step 1>
```

51436 IH-6. 6 

## Student Notes

Once you have chosen an operating system during the attended boot process, the system will begin to boot. This may take thirty seconds or more since the installation media is slower than your root disk. You will then see the familiar boot messages along with some messages that are not as familiar. Some interfaces will misidentify itself. The warning "Insufficient space for dump device ... " should also be ignored. A warning about not saving a complete core dump is not a problem because the root file system is mounted read only and is incapable of saving a core dump of any size. The warning "Unable to configure dump device ... " should also be ignored for the same reason.

After some time, approximately thirty seconds, the HP-UX kernel loaded on the installation media from the file /hp-ux begins running. This is a special version of the HP-UX kernel. After the operating system on the Installation media has booted, it automatically runs an interactive application program we refer to as the the `install` program. This program first presents a `Welcome Menu` with a prompt.

The "Welcome Menu" gives you an overview of what actions the installation will perform. You should press (Return) AFTER you have read it.

application to correct the assumptions about directories using the directory file information required by long file names.

- Programs (with no source code available) that were developed for or compiled on releases of UNIX that do not support long file names will be run on the system.
- Other systems in your organization run version of HP-UX that impose a 14 character limit on file name length. In this environment, you might want uniformity across the systems so that files can be moved between systems.

The Series 300/400 installation process does not provide this screen to the system administrator. Short file names are the default. However if you are on a Series 300/400 or on a Series 700 and decide to create your file system with short file name, a command exists, `/etc/convertfs`, that converts and existing file system supporting the short file names into one that supports long file names. Once a file system is converted to long file names, it cannot be restored to its original state.

### Tips on Using Menus:

- Arrow keys: Use to highlight menu items.
- `Select Item`: Use to select the highlighted menu item. (`Select Item` is one of several labels along the bottom of menus throughout the installation process. These labels correspond to function keys on the top row of your keyboard `f1`,`f2`, ... `f8`. The `Select Item` function key will always be `f4`.
- `Return`: Same as `Select Item` function key. Also use after entering values you have typed at the keyboard.



59

## IH-8. SLIDE: Destination Menu

33

### Destination Menu

```
HP-UX INSTALLATION UTILITY--DESTINATION MENU

    Select one of the following disks (name and system location)
    connected to your system to be the destination device for the
    installation.  Highlight the desired disk and press "Select Item."
    If the desired disk is not listed, make sure it is connected
    properly and turned on, then select the "Search Again" option.
    If your disk is STILL not recognized, you can use the
    "Other disk" option to manually enter the Disk address.

                  Select  Bus   Unit  Vol
 Disk             Code    Addr  Num   Num
----------------------------------------
HP 2213A          at 14    6     0     0
HP 7937           at 15    0     0     0
HP 7959S          at 14    5     0     0
HP 7935           at 15    1     0     0

Search Again
Other disk

[      ] [      ] [      ] [Select Item] [      ] [      ] [      ] [Exit Install]
```

51436 IH-8. 8 

## Student Notes

### The Destination Menu

Before seeing the Destination Menu, we will first get a message:

`-- Searching for possible destination disks.  Please Wait.  --`

After several seconds (10 to 25 seconds in some systems with lots of disks connected) we will then see the Destination Menu presented in the slide.

The Destination Menu lists all the hardware addresses of connected disks found by the boot ROM search. Choose the disk that matches the hardware address that you decided on earlier for the root disk. Use the arrow keys to highlight the item you want and then press `Select Item`.

On the Series 700, the choices will look a little different. A typical list of choices follows:

```
                    Slot   Bus  Unit
 Disk               Number Addr Num
---------------------------------------
QUANTUM PD210S   at  0      6    0
QUANTUM PD210S   at  0      5    0
HP 2213A         at  0      1    0
HP 2213A         at  0      0    0
```

If you do not see your destination disk listed on the menu, check to make sure that your disk is connected properly and is turned on. Then select "Search Again."

### Potential Disk Problems

The installation program should be able to find the disk that you expect it to find. If it does not, you might not have connected the disk properly. Loose cable connections are often the culprit, and sometimes the disk is on the wrong interface cable. Be sure the disk power is plugged in and the power switch is turned on! Remember that SCSI cables MUST be properly terminated. Check that the disk is properly connected, power is on and ready, and then select "Search Again."

If you still don't see your destination disk, select the "Other disk" option and type in the correct information.

Once you have chosen your destination disk, the install program checks for an HP-UX system already on the disk before you will be put into the "HP-UX INSTALLATION UTILITY—MAIN MENU." If an existing HP-UX system is found, the softkeys will be cleared, and a warning message will be given:

```
Warning: There appears to be an HP-UX system already on this
disk.
Press [Return] to continue.
```

This is normal, since the only acceptable response at this point is to press the Return key. The softkeys will be redrawn with the next screen.

If your goal is to overwrite this disk, simply press [Return] to continue, ignoring the warning. If you have selected the wrong destination, you will be given a chance in the next screen to select a different destination disk.



6c

## IH-9. SLIDE: Main Menu

34

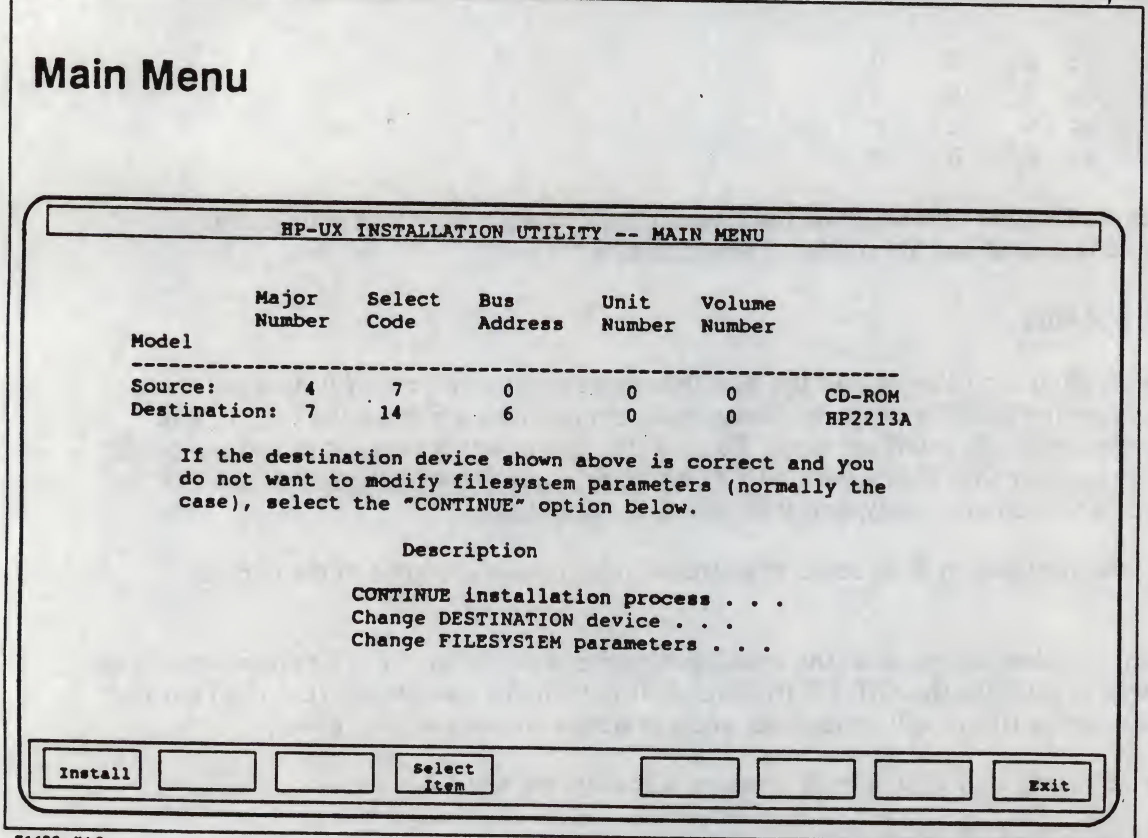


51436 IH-9. 9 

### Student Notes

After you select your destination device on the previous screen, you will see this menu. This menu lets you continue the installation, change your destination device, or change file system parameters.

The Series 700 hardware choices will differ slightly in format of hardware address.

| | Major Number | Slot Number | Bus Address | Unit Number | Model |
|---|---|---|---|---|---|
| Source: | 4 | 0 | 4 | 0 | Tape |
| Destination: | 7 | 0 | 0 | 0 | HP 2213A |

Usually you will select the choice to "Continue installation process".

The installation process initializes the file system using default parameters for the type of disk that you selected. The values for the default parameters are suitable for most HP-UX users. You should not need to change anything unless the default swap space is too small for your requirements.

If you are satisfied with the addresses of the source and destination devices, select "CONTINUE installation process." If you are not satisfied with your destination device, you have another chance to change it. Select "Change DESTINATION device" and you will be returned to the "Destination Menu." Re-enter your destination and then you will be returned to this menu.

If you select "Change FILESYSTEM parameters", you will be given the opportunity to change your file system and disk characteristics.



61

## IH-10. SLIDE: Filesystem Parameters Menu

35

### Filesystem Parameters Menu

```
HP-UX INSTALLATION UTILITY -- FILESYSTEM PARAMETERS MENU

The only parameter below that you may want to change is "Swap Size".
All others should be correct.

Swap space (512 Byte blocks): 86016  Swap space available on disk.
------------------------------

Block size:                    8192  File System Performance Params
Fragment size:                 1024
Rotational Delay (millisec):   4

Free Space Threshold (%):      10    File System Space Parameters
Density - bytes per inode:     2048
Cylinders per group:           16

1024 byte sectors per track:   16    Disk Geometry Parameters
Tracks per cylinder:           12

Interleave Factor:             1     Disk Hardware Parameters
RPMs of the disk:              3350

Highlight the parameter you wish to change. Press "Done" when finished.
           New Swap Space value >

                        Done
```

51436 IH-10. 10 

## Student Notes

**Note**

You should not change the values of any other file system parameters (besides swap space) unless you have expert knowledge of HP-UX and know why a parameter needs to be changed (and the value to which it must be changed). Only make changes if you have accurate and specific knowledge of the required change, do not guess. You might be directed to make this change by a "README FIRST" notice shipped with your installation documentation, for example. Failure to observe this precaution may render your system unusable.

You should enter the desired change (if any) and then press "Done". Changes are *not* usually required. The values shown here are for an HP 7959S SCSI drive selected as the destination root disk. The values you see may differ if you have another disk type.

**Note**

Be aware that the choices made here are **permanent** until the root disk is re-initialized, which implies destroying all the data on the disk. A complete system backup is critical to avoid data loss in that circumstance.

## IH-11. SLIDE: Swap Space Verification Menu

36

### Swap Space Verification Menu

```
Swap Space Verification

The current swap space is 98502 blocks (1 block = 512 bytes).
Is this correct? (Enter y or n) >>

Select Item                                    Exit Install
```

51436 IH-11. 11 

## Student Notes

Once you have selected "CONTINUE installation process" in the Installation Utility Main Menu (on the previous slide), you will be taken to a "Swap space verification" menu. This occurs even if you have changed the default swap space using the screen described in the notes.

No softkeys are displayed on this screen, since the only acceptable answer is "y" or "n" followed by pressing (Return).

If the swap space displayed is incorrect, answer "n", and respond to the prompt. You should input the number of 512- byte blocks of swap space you want to reserve on your system disk. This is the primary swap space. The amount you actually get will be rounded up to satisfy system constraints. You will again be asked to verify the amount.

Remember that it is very difficult to modify the swap space after you have installed your system, so make sure you are comfortable with the amount of swap space you choose. Once you have decided that the amount of swap indicated is correct, answer "y" followed by pressing (Return).

Remember that the series 700 uses 2MB from primary swap for boot up.

Your swap space must be large enough to hold all the processes that can be running at peak times. But if you have never had your system up and running, it is hard for you to estimate the swap size needed and to know which applications run at the same time. In general, the default value should be sufficient. Estimating the amount of swap you need can be difficult. See the Module SW for details.

However, if you are running very large applications or your system will be a diskless server, the default amount of swap may be too small. In these cases, you should consider increasing your swap space. Consult the "README FIRST" documents supplied with the application software for guidance in configuring swap space. The time to do that is BEFORE you reach this point, as was pointed out earlier in this module in the slide on "What it Means to Install HP-UX", in the section "Before You Start".

On the Verification Menu, the swap space is in blocks, where 1 block = 512 bytes. Therefore, you need to convert your swap space from megabytes to blocks before confirming the amount.

```
swap = swap_blocks x 512 bytes/block
```

so,

```
swap_blocks = swap / 512 bytes/block
```

An example using 50Mbytes of swap:

```
1 Mbyte = 1024*1024 bytes
```

so,

```
50 Mbytes = 50*1024*1024 bytes = 52,428,800 bytes
```

Recall,

```
swap_blocks = swap / 512 bytes/block
```

so,

```
swap_blocks = 52,428,800 bytes / 512 bytes/block = 102,400 blocks
```



63

## IH-12. SLIDE: Destination Device Verification Menu

37

### Destination Device Verification Menu

```
                    Destination Device Verification

      Destination Device: HP 2213A    at  14    6    0    0
      ------------------------------------------------------

  WARNING!  If you continue with the installation process you will
  destroy everything on your destination disk.  If you have information
  on your destination disk you with to save, you should not continue.

  Do you wish to continue? (Enter y or n) >>

                            Select                              Exit
                            Item                                Install
```

51436 IH-12. 12 

## Student Notes

Note that no soft keys are displayed in this screen, since the only acceptable response is "y" or "n". After entering "y," you will be given the option of initializing your destination disk. Then the core filesets of the HP-UX system will be loaded.

### Your Last Chance!

After you verify your swap space, the system will give you one last chance to change your mind before continuing. You will see a menu similar to that shown on the slide.

If you do not wish to continue, enter "n" and you will be returned to the Main Menu.

If you are satisfied with the destination and the swap space values, press "y." You will not have any other chance to change these values.

### Initialize Destination Disk

Before loading HP-UX onto your destination disk, you will be given the opportunity to initialize the disk. You must initialize if:

- your disk is known to be damaged, or
- you want a different interleave than what the disk currently has

Many of the SCSI disks require a very specific interleave factor for proper media initialization. If you are not certain that the interleave factor is correct for you destination disk, then you should not initialize your disk.

If you decide to initialize your disk, and respond with "y", it will be some time before more input is required. You will see a message like:

```
It will be about 20 minutes before any input is required.
Initializing destination disk.
```

The amount of time depends upon the type and size of your disk (intervals of time run from 10 or 20 minutes for high performance disks to 45 minutes for much slower disks). The initialization destroys all data previously existing on that disk, verifies the storage integrity, and marks any defective blocks so they will not be used.

### Create an HP-UX File System and Load HP-UX Core Filesets

Regardless of whether you initialized your disk or not, you will leave the "Initialization Menu," and enter the next step in the installation process. The next step will occur after initialization completes if you type "y". In this step an HP-UX file system is created on the disk. You will see a message similar to this:

```
It will be about ten minutes before any input is required.
Making file system on destination device.
```

After the file system is created for you, core filesets of HP-UX will automatically be loaded onto your root disk. This is a minimum set of files required to boot HP-UX from your disk and then run the update program. The update program will load all the rest of your operating system files and filesets. The core file sets contain a stripped down HP-UX kernel, a shell, and a minimum set of commands and utilities, and other tools for basic system administration tasks.

More filesets will be loaded in the rest of the installation process. For users who have ordered only the **HP-UX Runtime System,** the core filesets are nearly all they require to run their application(s).

It takes between 10 and 60 minutes to install the core filesets. After this completes, the system will reboot itself, this time from the destination device (your system disk). As the core filesets are being transferred, you will see messages similar to those below:

```
        HP-UX INSTALLATION UTILITY  -- EXECUTION TRACE

It will be about ten minutes before any input is required.
Making file system on destination device
Mounting destination device
Making directory: /dev
```



69

```
Making directory: /dev/dsk
Making directory: /dev/rdsk
        .
        .
        .
```

Then the installation utility program will clear the console screen, followed by these messages as the core filesets are unloaded onto the root disk:

```
     Unpacking tar(1) files
x ./etc/group, 142 bytes, 1 tape blocks
x ./etc/inst.update, 536576 bytes, 1048 tape blocks
x ./etc/passwd, 279 bytes, 1 tape blocks
        .
        .
        .
     Done unpacking files
        .
        .
        .
     Creating /etc/checklist
        .
        .
        .
```

A message about syncing the disks follows after the filesets have been loaded. The system will be **automatically** rebooted, this time from the minimal system just copied to the disk. You will see your console screen clear, and then start to display messages similar to the "Booting from Install Media" slide above.

## IH-13. SLIDE: Rebooting From System Disk

30

### Rebooting From System Disk

```
Copyright 1989,
CONSOLE is 98644 at sc 9,
1 port(s)
 . . .
Internal HP-IB Interface - system controller at select code 7
  Parallel poll interrupts enabled.
HP98644 RS-232C Serial Interface at select code 9
HP98265A (SCSI Interface) sync 2.67 MBytes/sec; parity
enabled at select code 14
HP98643 LAN/300 Link at select code 17 ignored; interface driver not present
real mem  = 4182016
 . . .
Initializing . . .

        Ensure that the install media unit has been removed and an
        update media unit is online and prepared for reading.

                --- Press "Return" to continue ---
```

51436 IH-13. 13 

## Student Notes

As the system reboots from the root disk, you will see now familiar messages on the screen. Next you will be prompted to remove the **INSTALL** medium and insert the first **UPDATE** medium into the installation device.

You must first unload the **INSTALL** medium. Next you should load the first **UPDATE** medium. Wait for the busy light to go off on the installation device. If this is a cartridge tape drive, the busy light should go out and remain OFF (it will blink on and off for several minutes). With a DDS format DAT, the tape will take about 25 seconds to load and the upper light will flash green. The upper light will be yellow for a write-protected tape. This is the case during an installation. (For a write-enabled tape, the upper light will be green.) With a CD-ROM the drive will be ready in about five seconds, the ready light will blink once very quickly and remain ON. When the installation device is ready, press Return. You will see the menu shown on the next slide.



65

**Warning**

**You MUST load the *UPDATE* media unit at this point and wait for it to become ready. You must first unload the *INSTALL* medium to do this. If you fail to swap the media as directed at this point when using the CD-ROM installation media, or continue by pressing (Return) before the CD-ROM drive is ready, you cause an unrecoverable error. You may continue for a few steps but will reach a point where you are unable to complete the installation process. If you are using tape, the program will allow you to recover from this error.**

## IH-14. SLIDE: "Update" Main Menu

39

### "Update" Main Menu

```
INSTALL                              Main Menu

     Highlight an item and then press "Return" or "Select Item."
     To refresh the screen, press CTRL-L.
     ------------------------------------------------------------

     Source: Tape Device              Destination: Local System
             /UPDATE_CDROM                          /
     ------------------------------------------------------------

          Select All Filesets on the Source Media
          Select Filesets for a Minimum System ->
          View/Select Partitions and Filesets ->

          Enter Codeword ->

          How to Use Install

 Help  |      | Shell | Select |      |      |      | Exit
       |      |       | Item   |      |      |      | Install
```

51436 IH-14. 14 

## Student Notes

The next phase in the installation process is to add more filesets to the existing HP-UX file system using a running version of HP-UX. This menu is actually the first menu of the update process. Details on update are covered in the Module UH.



66

## IH-15. SLIDE: Customizing the System

40

### Customizing the System

You will need to provide the following:

- Continent where your system will be located.
- Time zone
- Current date (day, month, and year).
- Current time of day (hours and minutes) using 24 hour clock
- The unique network host name assigned to your system.
- The IP (Internet Protocol) address assigned to your system

51436 IH-15. 15 

## Student Notes

The rebooting at the end of file loading is nearly the end of the **INSTALL** process. You must interact with a special program run only once at the first reboot. Completing that interaction represents the end of the installation process, but not the end of your responsibilities for the installation.

Eventually as the system reboots, you will see a `Console login:` prompt. Before the prompt appears you will be required to answer a series of questions to set up system customization files. Once you get the console login prompt, you must log in as `root` (no password is required) and perform additional postinstallation procedures (covered in the next module).

Following the reboot, you will see system startup messages on the console as before. The messages will be interrupted just after the `/etc/rc` script starts, with a series of questions that you must answer to customize your system for your locale and time of day.

Before you respond you will need to know the following data:

Continent where your system will be located.

Time zone
Current date (day, month, and year).
Current time of day (hours and minutes) using 24 hour clock
The unique network host name assigned to your system.
The IP (Internet Protocol) address assigned to your system

You may use the defaults for the network host name and the IP address if they are not known.

The messages you will see are listed below. Your first question will be:

```
/etc/rc:

################################################################################
The following procedure will allow you to set the time zone.

Select your continent from the following list:
              1. North and Central America
              2. South America
              3. Europe
              4. Africa
              5. Asia
              6. Australia, New Zealand

Enter the number corresponding to your continent (1-6) -> 1
```

The correct answer for the USA is 1. Next you will be asked:

```
################################################################################
Select your time zone from the following list:

1. Newfoundland Std. Time (NST3:30NDT) |  7. Mountain Standard Time (Arizona)
   Newfoundland Daylight Time          |     (No daylight savings time.)
                                       |
2. Atlantic Standard Time (AST4ADT)    |  8. Pacific Standard Time (PST8PDT)
   Atlantic Daylight Time              |     Pacific Daylight Time
                                       |
3. Eastern Standard Time (EST5EDT)     |  9. Yukon Standard Time (YST9YDT)
   Eastern Daylight Time               |     Yukon Daylight Time
                                       |
4. Eastern Standard Time (US:Indiana)  | 10. Aleutian Standard Time (AST10ADT)
   Central Daylight Time (EST5CDT)     |     Aleutian Daylight Time
                                       |
5. Central Standard Time (CST6CDT)     | 11. Hawaii
   Central Daylight Time               |
   Central American countries          | 12. Unlisted time zone
                                       |
6. Mountain Standard Time (MST7MDT)    | 13. Previous menu
   Mountain Daylight Time              |
```



67

```
Enter the number corresponding to your time zone (1-13) -> 5
```

The correct answer will obviously vary depending upon your site. Next you will be asked a series of questions about the date and time:

```
#################################################################################

You will be prompted for the date and time.  Please enter
all
values numerically, for example January is 1.  The values in
the parenthesis give the acceptable range of responses.

Please enter the month (1-12) -> 1
Please enter the day of the month (1-31) -> 25
Please enter the hour (using 24 hour time) (0-23) -> 15
Please enter the minute (0-59) -> 16
Please enter the last two digits of the year (70-99) -> 91

The date and time have been set to: Fri Jan 25 15:16:00 CST 1991
```

Your response of course varies, but should be relatively accurate. Next the system asks you for networking information. You should have already selected a host name and an IP Address for your system. You may use (Return) as a default.

```
For the system to operate  correctly  you must assign it a
system  name.  It should be 8  characters  or less and may
contain any alphabetic character, 0 through 9, -  or _
(hyphen/dash and underscore).
The  first   character   must  be   alphabetic.  Uppercase
alphabetic  characters  are  allowed but not recommended.
Entering nothing followed by return will enter the default
name of "unknown".


Enter the system name  -> training
```

If you provided an acceptable host name, you will then be prompted for a unique IP address. As you can see reasonable defaults are provided:

```
##################################################################

If you wish networking to operate correctly, you must assign the
system an Internet Protocol (IP) address.  This  address must be
assigned  by your local  network  administrator.  If you do know
your IP address you may press  return to use the default  value
(127.0.0.1).  An IP  address  should  consist  of  four numbers
separated  by periods.  The numbers  should be between 0 and 255
```

```
    inclusive.  For example: 255.32.0.10


Enter your Internet Protocol address -> 192.6.241.71
```

At this point You must simply wait for the system to complete the rebooting process. Your system will become very busy for a few minutes then complete the boot process.

At the end of the process your screen will have a number of

messages such as:

```
Starting up standalone system
Is the date Fri Jan 25 15:16:35 CST 1991 correct? (y or n, default: y) y
```

You may reply or let it time out and take the default of "y". This will be followed by more boot up messages. Many will indicate some form of error. They may be safely ignored at this time. You will see messages similar to this:

```
/etc/nettl: !: not found
ifconfig: socket: Protocol not supported
lanconfig: socket: Protocol not supported
ifconfig: socket: Protocol not supported
/bin/nodename: system feature not installed
Network Link started
System message logger started
    starting NFS networking
        NIS domain name not set
        /etc/portmap
    Network Information Service not started.
ARPA/Berkeley daemons started: inetd
Network NS Services started
Network management daemons started: socket: Protocol not
supported

NETWORKING started.
cron started
starting the ptydaemon
starting the vtdaemon
NOTE:  Files in /tmp:
total 236
crw-r--r--   1 root     other      4 0x070000 Nov  7 17:00 _chk_for_rombo
-rw-r--r--   1 root     other        765 Jan 25 15:19 nettlgen.log
-r-xr-xr-x   1 bin      bin         9096 Nov  6 14:54 rgb.800
-rw-r--r--   1 root     other      32480 Nov  7 23:24 rmfn.log
-rw-r--r--   1 root     other         24 Nov  7 23:26 update.cleanup
-rw-r--r--   1 root     other      75858 Nov  7 23:26 update.log
/etc/auditrc: This file must be edited to activate audit subsystem
Fri Jan 25 15:19:40 CST 1991

Console Login:
```



65

Finally at the console login, you may become the superuser by typing "root" (Return), and proceed to the post installation procedures in the next module.

## IH-16. LAB: Installing HP-UX

### Directions

With your lab partner, perform the following task.

1. Install HP-UX on your system or perform an installation on the demonstration system. Check with the instructor before performing this step. If a demonstration system is used, the instructor may ask a student to perform the installation with guidance from the rest of the class.



69

# Module CD — Device Files (Series 300/400/700)

## Objectives

Upon completion of this module, you will be able to:

- Describe how a device file is used.
- Differentiate between block and character I/O operations.
- Select an appropriate name for a device file using HP-UX naming conventions.
- List the device files needed by an HP-UX system.
- Describe major and minor numbers.
- Find major numbers for various types of interfaces.
- Find the components need to construct minor numbers.
- Create device files with `mknod`.
- Decipher the information displayed by the `ls -l` command when used on device files.

## IH-16. LAB: Installing HP-UX

### Directions

With your lab partner, perform the following task.

1. Install HP-UX on your system or perform an installation on the demonstration system. Check with the instructor before performing this step. If a demonstration system is used, the instructor may ask a student to perform the installation with guidance from the rest of the class.



69

# Module CD — Device Files (Series 300/400/700)

## Objectives

Upon completion of this module, you will be able to:

- Describe how a device file is used.
- Differentiate between block and character I/O operations.
- Select an appropriate name for a device file using HP-UX naming conventions.
- List the device files needed by an HP-UX system.
- Describe major and minor numbers.
- Find major numbers for various types of interfaces.
- Find the components need to construct minor numbers.
- Create device files with `mknod`.
- Decipher the information displayed by the `ls -l` command when used on device files.

## CD-1. SLIDE: Introduction To Device Files

41

### Introduction To Device Files

- Device files, or special files, allow the kernel to communicate with peripherals
- Each I/O device has one or more device file(s)
- `mknod` command creates device files

51436 CD-1. 1 

## Student Notes

## Student Notes

HP-UX communicates with peripheral devices such as tape drives, disk drives, printers, terminals, and modems through files called *device files* or *special files*. Because of the way *device files* are handled by HP-UX, I/O operations to peripheral devices are the same as I/O operations to ordinary disk files. This means that your programs do not have to distinguish between output destined for a file and output destined for a device.

Before HP-UX can communicate with a peripheral device, it must have a device file for that device. For example, each terminal has its own device file through which HP-UX writes data (which appears on the terminal screen) and reads data (typed by the user at the keyboard).



70

# Module CD — Device Files (Series 300/400/700)

A device file does not contain data as a regular file does. Instead, a device file specifies how HP-UX is to communicate with a device and where a device is located. Device files are created with either the `mknod(1m)` command or by using SAM. In HP-UX device files are stored in the `/dev` directory. Once created, they remain on the disk (in the `/dev` directory or in subdirectories below it) just as other files until they are explicitly removed.

**Note**

The term **device file** is equivalent with **special file**. Both of these terms are used in this module and mean the same thing. If you see a reference to **device special file** in other written material about HP-UX or UNIX, the writer is being conservative in his word choice and means device file or special file.

## CD-2. SLIDE: mknod - Creating Device Files

42

### mknod - Creating Device Files

mknod *name* c|b *major minor*

where:

| | |
|---|---|
| *name* | path name of device file to be created |
| c\|b | specifies character or block access to device |
| *major* | Kernel driver to be used to perform the data transfer |
| *minor* | Device location and driver specific information |

interface
+ "Scusi numeri"

51436 CD-2. 2 

## Student Notes

A device file gives the kernel access to information about a specific device. This information is essential to HP-UX when performing Input/Output operations with the device. HP-UX device access requires that a device file be created before any I/O operation can take place between it and the system.

Device files are created with the mknod command. Any time a file of any type is created a data structure that describes that file is created. This data structure is called the the information node or inode. In the case of a inode for a device file, the inode indicates that this file is a pointer to a device and not an ordinary disk file. Because of the similarities between the inodes of device files and regular files, HP-UX allows you to refer to devices and disk files interchangeably in your programs or commands. This is an important contribution of UNIX to ease of use.



71

`mknod` *name* `c|b` *major minor*

where:

| | |
|---|---|
| *name* | *name* is path name of device file to be created. |
| c\|b | The c or b indicates that the file is special file of type character or block. A device special file is either one or the other. |
| Major number | The device file's major number indicates which device driver to use. A driver is the part of the kernel software used to perform the actual data transfer between the device and the computer. |
| Minor number | The device file's minor number identifies to the kernel where it (the kernel) can find the device. Information such as which interface card handles this device and in the case of multiple devices located off of this interface card, which device is also identified. Additional information that controls the device behavior may also be encoded into the minor number. This information controls how the selected device will actually respond to the data transfer requests. |

Some devices will have two or more device files depending on how the driver is written. Some reasons for multiple device files include different behavior information or multiple device file names.

## CD-3. SLIDE: How HP-UX Device Files Are Used

43

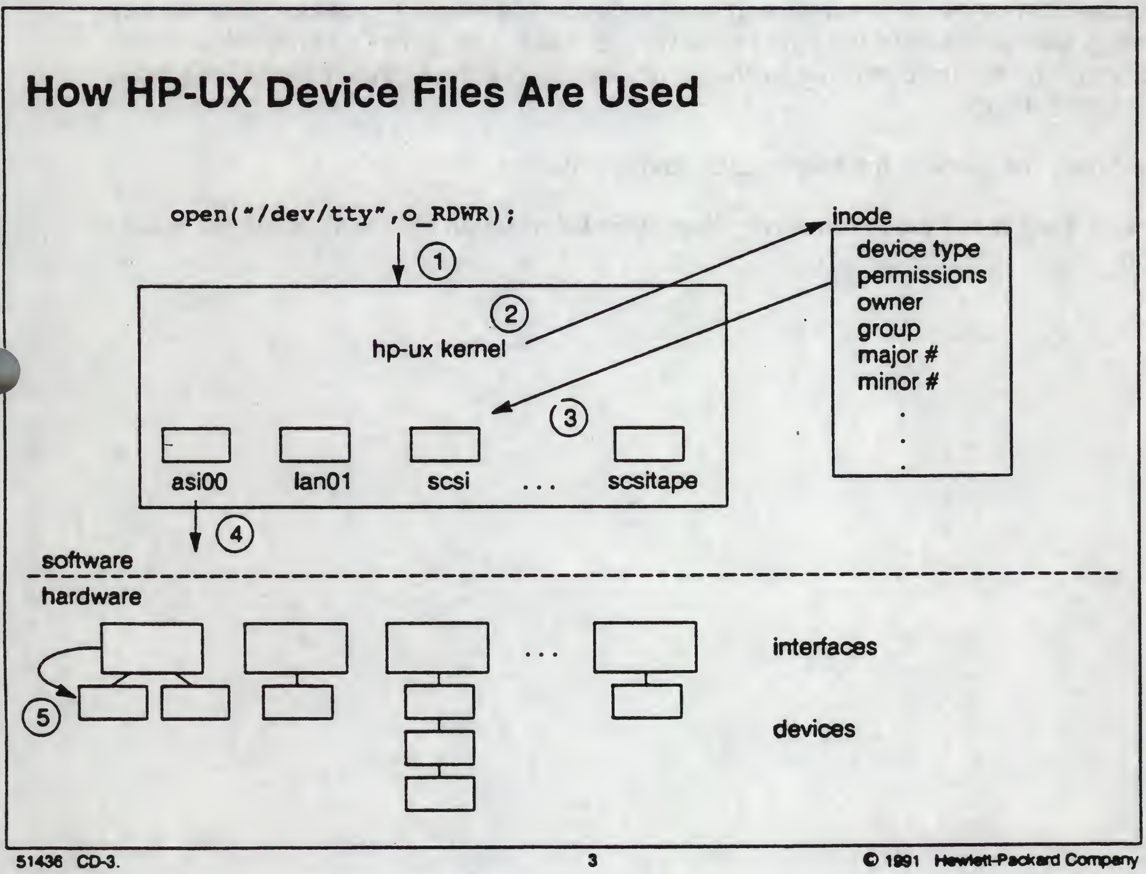


## Student Notes

Now that we generally know the components of a device file, we will follow the flow of an I/O operation to see the pivotal role the device file plays in HP-UX I/O.

Refer to the numbers on the slide, as we trace a "generic" I/O operation:

1. Your application requests an I/O operation on a specific file. A device file is named in this example. This request is made by a "system call". (There is a separate course covering system calls called "Programming with HP-UX System Calls").
2. The kernel gets the request for an I/O operation on a specific file by your application. The kernel then "opens" the file on behalf of your application. The kernel as part of the "open" discovers that it is a device file.



72

3. The kernel extracts the information about the device from the device file's inode and then calls the specified driver code (in the kernel) to carry out the request.
4. The driver accesses the interface specified at the hardware address specified in the minor number of the device file's inode.
5. The interface makes the I/O request to the device at the hardware address of the device. This part of the hardware address is also specified in the inode structure as was the interface's hardware address. Depending on the device, the device controller in the device may also need to direct the I/O request to the appropriate unit in the device.

   The I/O operation, a "read" or "write" for example, is carried out.

As you can see, the device file's inode holds some very important information that is necessary to make a successful I/O operation.

## CD-4. SLIDE: Block and Character Devices

44

### Block and Character Devices

Character devices

- I/O in character data streams
- Disks or tapes not holding a mountable file system
- Terminals
- Printers
- Modems
- All other devices

Block devices

- I/O in block units
- Disks or tape drives holding a mountable file system

51436 CD-4. 4 

## Student Notes

All I/O devices can be classified as either block or character devices.

**Block** device files transfer data in fixed sized blocks through system buffers, also called a kernel buffer cache to speed up I/O transfers. Devices that use block device files are random access storage devices such as disk drives and in some cases, tape drives.

**Character** device files typically transfer data in data streams without using any system buffer cache. This refers to any device that does not hold a mountable file system. The following devices use character device files: terminals, printers, plotters, digitizers, magnetic tape drives, cartridge tape drives, and, on occasion, disk mass storage devices. Character I/O is also called raw I/O transfer, and sometimes character devices are called raw devices.

# Module CD — Device Files (Series 300/400/700)

Recall that a character in HP-UX is the same as a byte. Thus I/O operations on a character device file moves data as a "stream of bytes" between the peripheral device and the system unit. In contrast, I/O operations on a block device file moves data as a "stream of fixed size blocks" between the peripheral device and the system unit.

A few devices are capable of I/O using either block or character device files. Such devices will usually have two device files: one for block access and one for character access. For example, a tape drive can use both block and character files. The block device file is used when you want to treat the tape as an HP-UX file system. This is usually done in emergency cases only since the tape unit has extremely slow random access to blocks. The character device file is used when you want I/O to occur in a stream. This is the normal mode for creating or restoring tape backups.

Hard disks should have both block and character device file entries since, in most cases, hard disks hold mountable file systems. Any cartridge or mag tape drive used for backup or recovery should have entries for both block and character device files. All other devices typically have only character device file entries.

## CD-5. SLIDE: /dev Directory

dd

45

mount

/tmp

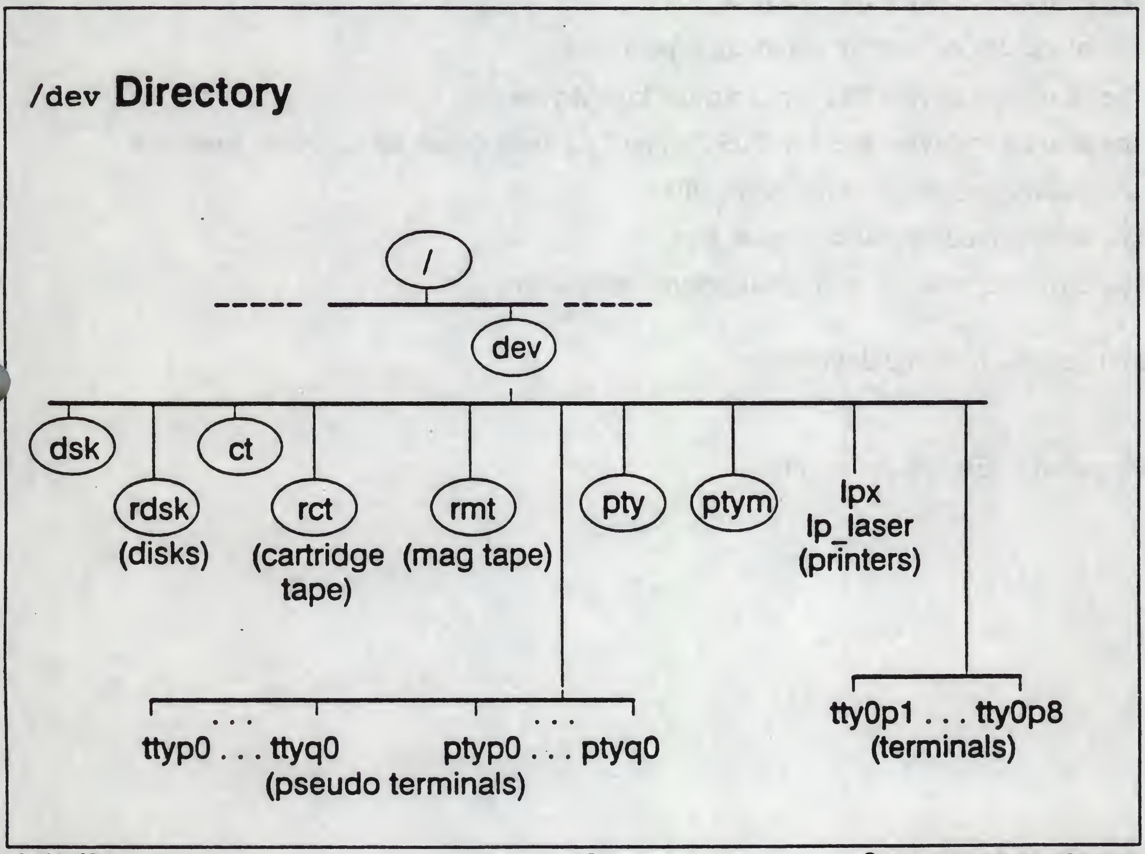


51436 CD-5. 5 

## Student Notes

By convention, device files are maintained in the directory **/dev** ("dev" for devices). Some device files are defined in the **/dev** directory itself while others are grouped in sub-directories under **/dev**. Device files that are defined in sub-directories are grouped by device type and by device file class.

There are subdirectories for device types include such groupings as cartridge tape, mag tape/DAT, and disk drives. Usually there are separate subdirectories for the block device files and character device files for a certain type of device, for example, disk drives. You will have two directories, one for block disk device files and another for character disk device files.



74

# Module CD — Device Files (Series 300/400/700)

The slide shows examples of some of the device files and sub-directories in /dev. Most of these sub-directories are created for us at installation time and are used to make it easier to locate and name device files.

| | |
|---|---|
| /dev/dsk | all the block device files for disk drives |
| /dev/rdsk | all the character device files for disk drives |
| /dev/ct | all the block device files for cartridge tape drives |
| /dev/rct | all the character device files for cartridge tape drives |
| /dev/rmt | all the character device files for DDS Format DAT tape drives and 1/2-inch mag tape |
| /dev/ptym | all the master pseudo terminal device files |
| /dev/pty | all the slave pseudo terminal device files |
| /dev/ttyXpY or /dev/ttyXY | all the terminal, modem, and serial printer device files |
| /dev/ptyp0, /dev/ptyp1, etc. | master pseudo terminal device files |
| /dev/ttyp0, /dev/ttyp1, etc. | slave pseudo terminal device files |

## CD-6. SLIDE: Device File Naming Conventions

46

### Device File Naming Conventions

N= busadres

6

| Type of Device | Type of Device File | Naming Convention |
|---|---|---|
| Disk (including Magneto-Optical) | Block | /dev/dsk/*N*s0 |
| | Character | /dev/rdsk/*N*s0 |
| CD-ROM | Character | /dev/rdsk/*N*s0 |
| | Block | /dev/dsk/*N*s0 |
| DAT and Mag Tapes | Character | /dev/rmt/*N*\|m\|h[n] |
| Cartridge tape | Block | /dev/ct/c*N* |
| | Character | /dev/rct/c*N* |
| Terminal | Character | /dev/tty*NN* or /dev/tty*N*p*M* |
| Modem (dial-in) | Character | /dev/ttyd*NN* or /dev/ttyd*N*p*M* |
| Modem(dial-out) | Character | /dev/cul*NN*, /dev/cua*NN* or /dev/cul*N*p*M*,/dev/cua*N*p*M* |
| Printer | Character | /dev/lp_*printer_name* |

density

uucp

call unit line

51436 CD-6. 6 

## Student Notes

The name of a device file can be chosen arbitrarily. However, we recommend that you follow these naming conventions. The naming conventions help make it easier to locate and identify a particular device file for a device. If you follow these conventions, the name of the device file should suggest something about its characteristics.

We have already seen that all device files are kept in the `/dev` directory, and that some device files are grouped in sub-directories under `/dev` by device type and by device file class (block or character). Now we will look at how to name the file itself.



75

## Disk

Within the /dev/dsk and /dev/rdsk directories, the following naming conventions are used to differentiate between the devices and disk sections.

/dev/[r]dsk/*N*s0

*N* is the HP-IB or SCSI bus address of the disk that is referenced by the special device file. This makes it relatively easy to figure out the device file associated with a particular disk drive and vice-versa.

The 0 after the s refers to the section number on the disk. Since the Series 300/400/700 computer does not support multiple sections per disk,so the section number is always 0.

This naming convention does not work very well if there are multiple SCSI or HP-IB cards installed with disk drives on more than one card since the same bus addresses may appear on different interfaces. You will need to invent a new naming convention in this case. You might want to encode other hardware address characteristics in the name, but generally follow the convention in order to still make it easy to infer those characteristics from the name.

You should be aware that the fsck command (which will be covered later) assumes that the same naming convention between the device files in /dev/dsk and /dev/rdsk directories.

Examples:

/dev/dsk/1s0

/dev/rdsk/1s0

/dev/dsk/6s0

The first example is the block device file for a disk at bus address 1. The bus address may refer to a SCSI bus address or the HP-IB bus address. There would be no difference in the file name, but there would be a difference in the minor number.

The second example is the character device file for the same disk at bus address 1.

The last example is the block special file for the disk at bus address 6. On a system that uses SCSI disks bus address 6 is usually the root disk or system disk that contains the bootable operating system. On a system using HP-IB disks the root disk is usually at bus address 0 - /dev/dsk/0s0.

---

**Note** To avoid potential serious problems with your system disk, there should not be a character special device file for the system (or root) disk since there is seldom a need to use it. If you do create such a raw device file, for use by dd for example, ensure the permissions and ownership is as follows:

```
crw-r-----   1 root    sys       47 0x201600 Mar 18 15:30 6s0
```

---

## CD-ROM Disk

`/dev/[r]dsk/Ns0`

CD-ROM uses the same device file naming convention as disk drives. *N* is the SCSI or HP-IB bus address of the CD-ROM drive.

HPIB CD-ROM drives should be on a separate HP-IB interface card from the system disk drive. This will usually yield improved performance for the disk drive. It is not necessary to place SCSI CD-ROM drives on a separate SCSI interface card from SCSI disks.

Examples:

```
/dev/dsk/1s0
/dev/rdsk/1s0

/dev/dsk/4s0
/dev/rdsk/4s0
```

The first two examples are the block and character special device file for a CD-ROM disk drive at SCSI or HP-IB bus address 1. The third and fourth examples are the block and character special device file for a CD-ROM disk drive at SCSI or HP-IB bus address 4.

You may want to create a device file name that is more descriptive to the user. This can be done with the `ln` (link) command or by making another device file with `mknod`. If you do this it is still a good idea to keep an entry for the CD-ROM drive in the /dev/rdsk and /dev/dsk directories to alert you when you are adding other mass storage. You might consider, for example:

```
# ln /dev/dsk/4s0 /dev/cdrom
```

## DDS Format DAT and Mag Tapes

`/dev/[r]mt/Nl|m|h[n]`

*N* is the SCSI or HP-IB address of the Digital Data Storage (DDS) Format Digital Audio Tape (DAT) or 1/2-inch mag tape drive.

The suffix letter `l|m|h` specify the density at which the tape is to be read or written.

`l` indicates low density (800 bpi)

`m` indicates medium density (1600 bpi)

`h` indicates high density (6250 bpi)

A device file ending in an `n` indicates that the tape will not be rewound or repositioned in any way when the device file is closed. This is usually referred to as setting the tape drive for "no rewind on close."

Examples:

```
/dev/rmt/0m

/dev/rmt/0hn
```



26

The first example is the character special file for reading or writing to a device at bus address 0. This device is either a mag tape at medium density (1600 bpi) or DDS Format DAT tape. The second example is the character special file for reading or writing a 1/2-inch mag tape at high density (6250 bpi) with no rewind on close also at bus address 0.

In the case of mag tape, the two device files in the examples might both refer to the same mag tape drive, since most mag tape drives can read and write at multiple densities. One device file is used for high density tapes and the other for medium density, but both are never used at the same time!

When possible, you should have your HP-IB tape drives (mag tape or DDS Format DAT and disk drives on separate HP-IB interface cards. This will usually yield improved performance of the disk drives. You may have an HP-IB tape drive on the same HP-IB interface as a disk if the tape drive is being used very infrequently. In that case you should consider keeping the tape drive powered off until it is needed to avoid interference with the HP-IB disks.

It is not necessary to place SCSI tape drives on a separate SCSI interface card from SCSI disks.

## Magneto-Optical Disk

```
/dev/[r]dsk/mo
/dev/[r]dsk/Ns0
```

*N* is the SCSI bus address of the Magneto-Optical disk drive (MO disk). We have two possible choices for naming device files for a MO drive. One choice is to retain the same device file naming convention as disk drives. Another possibility is to assign a descriptive name to the device files. Since these devices have lower performance, and have removable media, the descriptive name helps distinguish the MO drive from other disks, avoiding potential problems.

Examples:

```
/dev/dsk/mo
/dev/rdsk/mo

/dev/dsk/2s0
/dev/rdsk/2s0
```

The first two examples are the block and character special device file for a MO disk drive. There is nothing that suggests the SCSI bus address when using the more descriptive device file name.

The third and fourth examples are the block and character special device file for the same device using the same device file naming convention as disk drives. The last two device file names suggest that the SCSI bus address is 2 for the MO disk drive. We may choose to have both names linked together.

It is not necessary to place SCSI MO drives on a separate bus from SCSI disks.

## Cartridge Tape

```
/dev/[r]ct/cN
```

*N* is the HP-IB address of the cartridge tape drive. When possible, you should have your cartridge tape drives and disk drives on separate HP-IB interface cards. This will usually yield improved performance of

the disk drives. You may have a cartridge tape on the same HP-IB bus as a disk if the tape unit is being used very infrequently. In that case you should consider keeping the tape drive powered off until it is needed to be used.

Examples:

```
/dev/rct/c1

/dev/rct/c2
```

The first example is the character special device file for a cartridge tape drive at HP-IB address 1. The second example is the character special device file for a cartridge tape drive at HP-IB address 2.

## Terminal, Modem, and Serial Printer

Terminal, modem, and serial printer special files have no subdirectory of their own and are kept in /dev directory. They follow the naming conventions listed below.

| | |
|---|---|
| /dev/tty*NN* or /dev/tty*N*p*M* | (direct connect, hard wired) |
| /dev/ttyd*NN* or /dev/ttyd*N*p*M* | (dial-in modem) |
| /dev/cul*NN* or /dev/cul*N*p*M* | (dial-out modem) |
| /dev/cua*NN* or /dev/cua*N*p*M* | (dial-out modem) (Both are required for modems) |
| /dev/lp*N* | (serial printer) |
| /dev/lp*printer-name* | (serial printer) |

*NN* is an arbitrary, unique number that serves as an identifier for an RS-232 port. By convention, ports are usually numbered sequentially starting at zero (00 for *NN*, 0 for *N*). *N*p*N* is usually employed for serial port that are on mux cards. The *N* is indicates which mux card, and *M* indicates which port on the mux.

Modems require three special device files and all three should use the same number (*NN* or *N*p*M*) to indicate they refer to the same modem. Two of the conventional device file names are for the autodialer and dial-out line units, which were separate devices originally. Although modern modems have integrated autodialers, the UNIX convention requires two otherwise identical device files.

Printers can be numbered starting at 0 or 1. Some administrators choose to name printers /dev/lp*printer-name*, where *printer-name* is the model number of the printer (for example, /dev/lp2235 or /dev/lp2686).

Examples:

```
/dev/tty00
```

is the special file for the first RS-232 terminal port created on the system.

```
/dev/ttyd00   (tty dial-in port 00)
```



37

is the device file for a dial-in terminal at port 0.

```
/dev/cul00   (calling unit, line 00)
/dev/cua00   (calling unit, autodialer 00)
```

are special files for a dial-out modem at port 0.

```
/dev/lp2235
```

is the special file for a serial printer (in this case a model HP 2235A).

**Note** If you use SAM to add terminals or modems, it follows the naming conventions described in this section. If you use SAM to add printers, it follows a slightly different convention, using an embedded underscore in the device file name.

For example if SAM added the printer called "laser" it would create a device file called "lp_laser" for you.

## CD-7. SLIDE: Listing Special Files

47

### Listing Special Files

ll
ls -al

```
$ ll /dev/console /dev/*dsk

crw--w--w-    1 rbw       timber     0 0x000000 Apr 16 12:56 /dev/console

/dev/dsk:
total 0
brw-r--r--    1 root      sys        7 0x410200 Apr 16 09:08 0s0
brw-r-----    1 root      sys        7 0x201600 Jan 14 16:41 6s0

/dev/rdsk:
total 0
crw-r--r--    1 root      sys       47 0x410200 Apr 16 09:08 0s0
crw-r-----    1 root      sys       47 0x201600 Mar 18 15:30 6s0
```

disk
Busadress

major
Select code

chmod
+x
-x
gid = RWX

51436 CD-7. 7 

## Student Notes

Some very important information can be obtained by doing a simple ll command on special files.

The first character of each line identifies the type of device file. A b denotes a block device whereas a c denotes a character device. The major and minor numbers appear immediately before the date. Recall that the major number directs the kernel to appropriate driver for this device. The minor yields information about location and device characteristics.



78

# Module CD — Device Files (Series 300/400/700)

**Note**

The permissions on the `/dev` directory should be 755. The owner and group should be root.

The table below lists the device files needed by the HP- UX system.

**Table CD-1.**

| | |
|---|---|
| syscon, systty, console | HP-UX uses these files to access the system console |
| kmem | Virtual memory uses this file to access the kernel |
| mem | Physical memory uses this file to access the kernel |
| null | This is a "bit bucket," used for unwanted output |
| tty | HP-UX uses this file to access a user terminal |

**Note**

Do not remove any of these files from the HP-UX system. All these device files (and more) are created for you at installation time.

The file `/dev/tty` is often confused with an ordinary "tty" device file. It is used with a special kernel driver to remap references to `/dev/tty` to the "real" tty port that an application is using. Applications are usually started from a session which has a "real" tty device (a terminal) associated with it, this is the device file to which references to `/dev/tty` are re-mapped by the driver (number 2) associated with `/dev/tty`. This useful and important device file should **never** be removed when deleting "tty" entries.

## CD-8. SLIDE: lsdev and Major Numbers

48

### lsdev and Major Numbers

drivers

Ser Karte RS232

```
# lsdev

    Character      Block        Driver
        0            -1         /dev/console
        1            -1         HP98628, HP98626, and HP98642
        2            -1         /dev/tty
        3            -1         /dev/mem, /dev/kmem, and /dev/null
       16            -1         Master pty
       17            -1         Slave pty
       23            -1         raw 8042 HIL
       24            -1         HIL
       25            -1         HIL cooked keyboards
       47             7         SCSI disk
       52            -1         IEEE 802 device     lan
       54            -1         SCSI tape
```

Bet geen driver

51436 CD-8. 8 

## Student Notes

The **lsdev** command lists device drivers in the system. The first two columns document character and block major numbers respectively. A -1 in either column means that a major number does not exist for that type of driver. The third column is text that describes the driver's use.

On the Series 700, the actual list of drivers in the kernel is listed. On the Series 300/400, an internal list, not the actual drivers in the kernel, is listed

Unfortunately, the wording in the column that describes the driver's use is not always as clear as a new system administrator may want. For example, the driver for serial devices such as terminals, printers, and modems is character driver number 1; not character driver number 2. HP928628, HP98626, and HP98642 were the original product numbers for HP serial cards. Character driver number 2 is used for



79

the device file `/dev/tty`. This special purpose device file refers to a process' control terminal. `/dev/tty` is useful for a process which must do I/O to the terminal in the case that `stdin` and or `stdout` has been redirected.

Other entries from `lsdev` output include:

**Table CD-3.**

| Character | Block | Driver | |
|---|---|---|---|
| 4 | 0 | CS80 disk & tape | HP-IB disk and cartridge tape |
| 7 | -1 | HP-UX printer | HP-IB printer |
| 8 | -1 | /dev/swap | |
| 9 | -1 | HP7974/HP7978 nine-track magnetic tape | |
| 11 | 2 | AMIGO disk | HP9121S/D |
| 21 | -1 | HPIB DIL | |
| 22 | -1 | GPIO DIL | |
| 55 | 10 | optical autochanger | Library System |

## CD-9. SLIDE: /etc/master and Major Numbers

49

### /etc/master and Major Numbers

```
*
* name          handle          type    mask    block   char
*
scsi            s2disk          3       3FB     7       47
scsitape        s2tape          1       FA      -1      54
parallel        CentIf          10      100     -1  -1
cs80            cs80            3       7FB     0       4
                        .
                        .
                        .
*
* The following entries form the alias table.
* field 1: product #    field 2: driver name
*
2213A           scsi
Quantum_210S    scsi
210S            scsi
s2tape          scsitape
s2disk          scsi
CentIf          parallel
HP_7937         cs80
                        .
                        .
                        .
```

tape / Bus

tape

Com Set

51436 CD-9. 9 

## Student Notes

Another source of information to use when searching for a major number is **/etc/master**. This file is used at system configuration time. It contains, among other things, device information. One of its nice features is the fact that actual product numbers are listed in this file.

A good example of this is the hp7937 HP-IB disk drive. **lsdev** would tell you:
4 0 CS80 disk
If you didn't know that a hp7937 used CS80 protocol, you might not figure out that the block and character major numbers are 4 and 0 for a hp7937. However, the **/etc/master** file actually lists the product number HP_7937's driver as CS80. You could then find the major numbers for the the CS80 driver.

## Module CD — Device Files (Series 300/400/700)

All supported device drivers should be listed in **/etc/master**.

## CD-10. SLIDE: Series 700 Minor Number

### Series 700 Minor Number

The minor number format is 0x*SSFDDD* where:

| | |
|---|---|
| 0x | Hexadecimal notation |
| *S* | System bus module number |
| *S* | EISA slot number |
| *F* | Function number |
| *DDD* | Driver specific information |

51436 CD-10. 10 

## Student Notes

The minor number format is 0x*SSFDDD* where:

0x Hexadecimal notation

*S* System bus module number
Recall the Series 700 Bus Architecture from Module HW that SBM numbers can have the following decimal values:

| | |
|---|---|
| 1 | Graphics |
| 2 | Core I/O board |
| 4 | EISA Bus Adapter |



81

| | |
|---|---|
| 8 | Processor |
| 9 | Memory |

The SBM number is 4 bits and as such is a single hexadecimal digit.

*S* EISA slot number
EISA slot number is only used when for cards located off of the EISA bus. The SBM is 4 for the EISA bus adapter. The EISA slot number is 4 bits and as such is a single hexadecimal digit.

*F* Function number
The function number is used to specify a particular function if the interface card is a multifunction card. For the core input/output interface function numbers are:

| | |
|---|---|
| 1 | SCSI |
| 2 | LAN |
| 3 | HIL |
| 4 | Serial Port 1 (labeled Port A) |
| 5 | Serial Port 2 (labeled Port B) |
| 6 | Parallel |

The core input/output's SBM is 2. The function number is 4 bits and as such is a single hexadecimal digit.

*DDD* Driver specific information
This information is 12 bits (three hexadecimal digits) and its format is different depending on the driver that is used to communicate with the device.

## Series 700 Bus Architecture

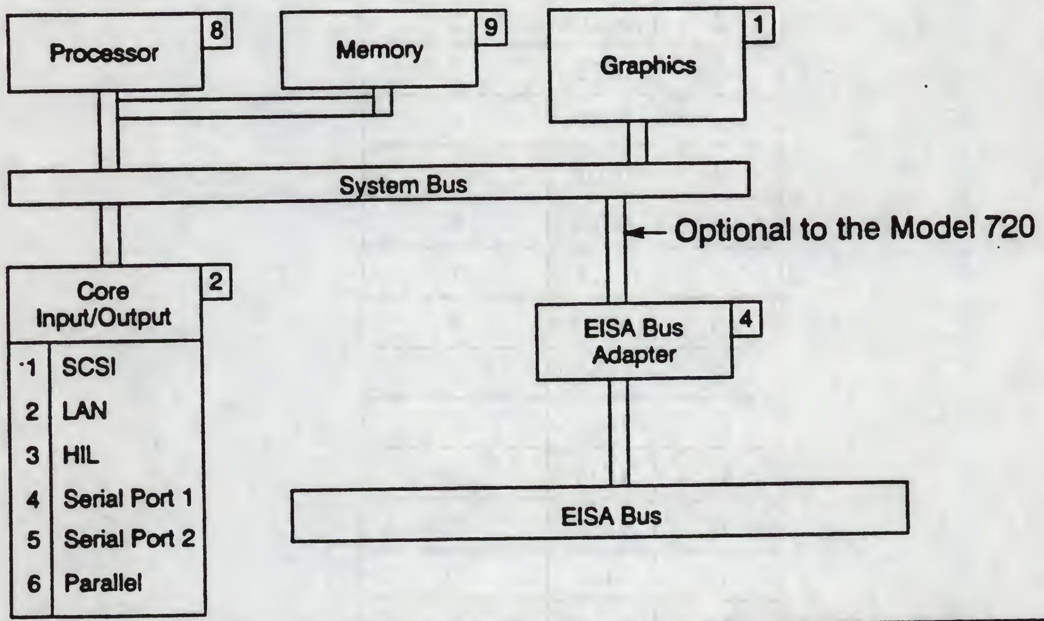


51436 HW-13. 3 

The following table provides a mapping of decimal to binary to hexadecimal values.

**Table CD-4. Decimal, Binary, Hexadecimal Equivalents**

| Decimal | Binary | Hexadecimal |
|---|---|---|
| 1 | 0001 | 1 |
| 2 | 0010 | 2 |
| 3 | 0011 | 3 |
| 4 | 0100 | 4 |
| 5 | 0101 | 5 |
| 6 | 0110 | 6 |
| 7 | 0111 | 7 |
| 8 | 1000 | 8 |
| 9 | 1001 | 9 |
| 10 | 1010 | A |
| 11 | 1011 | B |
| 12 | 1100 | C |
| 13 | 1101 | D |
| 14 | 1110 | E |
| 15 | 1111 | F |

## CD-11. SLIDE: Series 300 Minor Number

### Series 300 Minor Number

The minor number format is 0x*ScDDDD* where:

| | |
|---|---|
| 0x | Hexadecimal notation |
| *Sc* | Select code |
| *DDDD* | Driver specific information |

51436 CD-11. 11 

## Student Notes

The minor number format is 0x*ScDDDD* where:

0x Hexadecimal notation

*Sc* Select code
Recall the Series 300/400 hardware discussion from Module HW. Select codes are hardware addresses associated with a card. Some multi-function cards have more than one select code. Each select code must be unique to the system. Select code are set prior to shipment to the customer, but may be changed by the customer by changing dip switches located on the interface board. On boot-up the Series 300/400 systems report select codes of their interface boards.



83

*DDDD* Driver specific information
This information is 16 bits (four hexadecimal digits) and its format is different depending on the driver that is used to communicate with the device.

## CD-12. SLIDE: Minor Number For Mass Storage Devices

fd = Inode

50

### Minor Number For Mass Storage Devices

400

The minor number format is 0x*SSFBUV* or 0x*ScBaUV* where:

pag HW22

| | |
|---|---|
| *B* | SCSI bus address number on a Series 700 |
| *Ba* | SCSI or HP-IB bus address number on a Series 300/400 |
| *U* | Unit number |
| *V* | Volume number |

00

Series 700 Examples:

```
mknod /dev/rdsk/3s0 c 47 0x201300
mknod /dev/dsk/3s0 b 7 0x201300

mknod /dev/rdsk/3s0 c 47 0x420300
mknod /dev/dsk/3s0 b 7 0x420300
```

1.

Series 300/400 Examples:

major minor

```
mknod /dev/rdsk/6s0 c 4 0x0e0600
mknod /dev/dsk/6s0 b 0 0x0e0600

mknod /dev/rdsk/3s0 c 47 0x0e0300
mknod /dev/dsk/3s0 b 7 0x0e0300
```

51436 CD-12. 12 

## Student Notes

Mass storage devices include all disks, including magneto-optical and CD-ROM, and cartridge tape devices. The minor number format is 0x*SSFBUV* or 0x*ScBaUV* where:

*B* — SCSI bus address number on a Series 700
SCSI bus addresses are usually set on the device by dip switches or thumbwheels. The bus address is four bits or single hexadecimal digit.

*Ba* — SCSI or HP-IB bus address number on a Series 300/400
SCSI and HP-IB bus addresses are usually set on the device by dip switches or thumbwheels. The bus address is eight bits or two hexadecimal digits.

*U* — Unit number
The unit number is typically 0. Exceptions include older devices that have two devices



84

in the same physical box at the same bus address. Examples of these devices are an hp7946 which houses both a disk and a built-in cartridge tape drive and a hp9122 with has two floppy drives. Also, a hp7907 which is a disk with a removable platter has more than one unit number.

*V* Volume number
The volume number is always 0.

## Series 700 Notes

```
mknod /dev/rdsk/3s0 c 47 0x201300
mknod /dev/dsk/3s0 b 7 0x201300
```

These mknod commands create device files for a mass storage device at core I/O at SBM 2, function number1, SCSI bus address 3. If the mass storage device is on the core I/O board, the SBM is 2 and the function number is 1. The slot number is not used.

```
mknod /dev/rdsk/3s0 c 47 0x420300
mknod /dev/dsk/3s0 b 7 0x420300
```

These mknod commands create device files for a mass storage device at EISA adapter at SBM 4, EISA slot number 2, SCSI bus address 3 If the mass storage device is off of the EISA adapter board, the SBM is 4 and the slot number will vary depending on where the SCSI card is plugged in. The function number is not used.

## Series 300/400 Notes

```
mknod /dev/rdsk/6s0 c 4 0x0e0600
mknod /dev/dsk/6s0 b 0 0x0e0600
```

These mknod commands create device files for a HP-IB mass storage device at select code 14, bus address 6, unit 0.

```
mknod /dev/rdsk/3s0 c 47 0x0e0300
mknod /dev/dsk/3s0 b 7 0x0e0300
```

These mknod commands create device files for a SCSI mass storage device at select code 14, bus address 3, unit 0.

If the mass storage device is on high speed HP-IB or SCSI, the select code is 14. If the mass storage device is on standard speed HP-IB, the select code is 7.

## CD-13. SLIDE: Minor Number for DDS Format DAT and Mag Tape

51

### Minor Number for DDS Format DAT and Mag Tape

The minor number format is 0x*SSFBDO* or 0x*ScBaDO* where:

*D* density and compression

*O* operation characteristics

**Tape Density bits 6 and 7**

| bit 7 | bit6 | Density | Device Name |
|---|---|---|---|
| 0 | 0 | 800 bpi | /dev/rmt/0l |
| 0 | 1 | 1600 bpi | /dev/rmt/0m |
| 1 | 0 | 6250 bpi | /dev/rmt/0h |
| 1 | 1 | compressed | /dev/rmt/0c |

**Tape Operation bits 0-1**

| bit 1 | bit 0 | style of close | rewind on close | Device name (for 1600bpi) |
|---|---|---|---|---|
| 0 | 0 | AT&T | yes | /dev/rmt/0m |
| 0 | 1 | AT&T | no | /dev/rmt/0mn |
| 1 | 0 | Berkeley | yes | /dev/rmt/0m |
| 1 | 1 | Berkeley | no | /dev/rmt/0mn |

51436 CD-13. 13 

## Student Notes

The minor number format is 0x*SSFBDO* or 0x*ScBaDO* where:

*D* density and compression
Density is used to select densities of 6250, 1600, or 800 bpi. This corresponds to the h, m, and l flags in the recommended device name. With DDS Format DAT devices (digital audio tape), density designations are not used. On tape drives that support data compression, selecting the device file with c in the recommended device name causes the data to be written or read in compressed mode.

*O* operation characteristics
Rewind on close: The tape is automatically rewound upon close. When a rewind is not



85

desired, the device name with the **n** flag should be used. The *O* option is used with both mag tape and DDS format DAT.

AT&T-style vs Berkeley-style: When a file oepn for reading only is closed and the no-rewind bit is not set, the tape is rewound. If the no-rewind bit is set, the behavior depends on the style mode. For AT&T style devices, the tape is positioned after the EOF following the data just read. For Berkeley-style devices, the tapes is not repositioned in any way.

## Series 700 Example:

`mknod /dev/rmt/0m c 54 0x201300`
This `mknod` command creates a device file for a DDS Format DAT tape on the core I/O board, function number 1 for SCSI, and bus address 3.

## Series 300/400 Examples:

`mknod /dev/rmt/0m c 54 0x0e0300`
This `mknod` command creates a device file for a SCSI DDS Format DAT tape on select code 14 and bus address 3.

`mknod /dev/rmt/0m c 9 0x070300`
This `mknod` command creates a device file for a HP-IB mag tape on select code 7 and bus address 3.

See `mt(7)` for details.

## CD-14. SLIDE: Minor Number For Serial Devices: Terminals, Modems and Printers

52

### The Minor Number For Serial Devices:

### Terminals, Modems and Printers

The minor number is 0x*SSF00A* or 0x*ScPa0A* where:

*F* Function number for Series 700

*Pa* Port address for Series 400/300

*0* Always 0

*A* Specifies the access type

0 Simple protocol (U.S.) dial-in
1 Simple protocol (U.S.) dial-out
2 CCITT protocol (Europe) dial-in
3 CCITT protocol (Europe) dial-out
4 direct connect

modems

Series 700 Example:

```
mknod /dev/ttyd0p1 c 1 0x205000
mknod /dev/cul0p1 c 1 0x205001
mknod /dev/cua0p1 c 1 0x205001
```

Series 400/300 Example:

```
mknod /dev/ttyd00 c 1 0x0d0000
mknod /dev/cul00 c 1 0x0d0001
mknod /dev/cua00 c 1 0x0d0001
```

selly

51436 CD-14. 14 

### Student Notes

The minor number is 0x*SSF00A* or 0x*ScPa0A* where:

*F* Function number for Series 700 This is the same function number that we have seen before. For the serial ports on the core input/output interface, 4 and 5 are the function numbers.

*Pa* Port address for Series 300/400

Some serial cards on a Series 300/400 computer support one or more serial ports. You will have to create at least one device file for each port. If the interface card has only one possible port, then the port number is zero. If the device supports multiple ports, as in a



86

serial Multiplexer interface, the port number is labeled on the interface. Port numbers are always numbered from zero.

The HP98642A four port MUX supports ports 0, 1, 2, and 3. The HP98638A eight port MUX appears to be two adjacent 4 port MUXs and as such requires two adjacent select codes. It is connected to an active distribution panel where there are two sets of four ports, each supports ports 0, 1, 2, and 3. If the K2292 adapter is used on a series 400, there are three separate select codes (which are 5,6,9), and only one port (number 0) on each select code.

*0* Always 0

*A* Specifies the access type

| | |
|---|---|
| 0 | Simple protocol (U.S.) dial-in |
| 1 | Simple protocol (U.S.) dial-out |
| 2 | CCITT protocol (Europe) dial-in |
| 3 | CCITT protocol (Europe) dial-out |
| 4 | direct connect |

### Device Files for Modem Support

A modem requires three device files: two for the dial-out modem and one for the dial-in terminal. The device file naming conventions for a modem are shown below.

- `/dev/cul`*NN* is the dial-out modem (used once the connection is established).
- `/dev/cua`*NN* is auto-dial unit (used to establish the connection).
- `/dev/ttyd`*NN* is the dial-in terminal.

It is common to create one of the device files using a hard link between cul*NN* and cua*NN*. For example:

```
# ln /dev/cul00 /dev/cua00
```

### Device Files for Direct Connect

Serial devices that connect directly to your system, such as terminals or serial printers, only require one device file. These are often called "hardwired" terminal ports, since there is usually a physical wire that runs from the system to the terminal with no modem in between.

The table below shows what the minor numbers would be for a dial-in, dial-out, or a direct connect device on the second serial port on the core I/O board.

### Series 700 Example

```
mknod /dev/ttyd0p1 c 1 0x205000
mknod /dev/cul0p1 c 1 0x205001
mknod /dev/cua0p1 c 1 0x205001
```

These `mknod` commands create devices files for modem on the core I/O board at SBM 2, function 5. Function 5 on the core I/O board is the second serial port.

**Series 300/400 Example**

```
mknod /dev/ttyd00 c 1 0x0d0000
mknod /dev/cul00 c 1 0x0d0001
mknod /dev/cua00 c 1 0x0d0001
```

These mknod commands create devices files for modem at select code 13, port 0.



87

## CD-15. SLIDE: Minor Number For Parallel Devices

53

### Minor Number For Parallel Devices

The minor number format is 0x*SSF00A* or 0x*Sc000A* where:

*0* Always 0

*A* Mode

Series 700 Example:

```
mknod /dev/lp c 11 0x206002
```

Series 400 Example:

```
mknod /dev/lp c 11 0x0c0002
```

51436 CD-15. 15 

## Student Notes

The minor number format is 0x*SSF00A* or 0x*Sc000A* where:

*0* Always 0

*A* Mode
The only supported modes is 2. This value is four bits, a single hexadecimal digit.

Series 700 Example:

```
mknod /dev/lp c 11 0x206002
```

This mknod command creates a device file for a device on the core I/O board at SBM 2. The parallel function number is 6. The mode is 2.

Series 400 Example:

```
mknod /dev/lp c 11 0x0c0002
```

This mknod command creates a device file for a device at select code 12. The mode is 2.

Series 300 Example:

```
mknod /dev/lp c 11 0x170002
```

This mknod command creates a device file for a device at select code 23. The mode is 2.

## CD-16. SLIDE: Review of Creating Device Files

54

### Review

To create device files do the following:

1. Choose a name, avoid conflicts with existing files
2. Decide if device is character or block (or both)
3. Find the major number(s) by:
   - Looking in `/etc/master`, find aliases, driver numbers

     or

   - Executing `lsdev` command to find driver numbers
4. Construct minor number after deterimining device location and characteristics
5. Use the `mknod` command to create the device file

51436 CD-16. 16 

### Student Notes

For more information, see *mknod(1m)* and *mkdev(1m)* Also consult the script `mkdev` created by `config -a` for a shell script with numerous comments.

## CD-17. SELF-STUDY QUESTIONS (Series 300/400/700)

### Directions

Complete the following questions on your own.

1. What is a device file used for?

2. What distinguishes a device file from regular or directory type files?

c
b
d

bij LL

3. What is the difference between a block and a character device?

4. In which directories are device files defined?

5. What are major and minor numbers?

driver nummer

Hardware

6. Why do you need both block and character device files for disks?

nieuwe disks

on filesystem aanmaken



89

## CD-18. LAB: Device Files (Series 300/400)

### Directions

For this exercise it is *not* necessary to actually make the required device files, as there may be conflicts with existing files on your training system. Your instructor will indicate if any of the device files are to be created on your training systems.

Write out the command you would use to create device files for the following devices:

1. A terminal connected to the built-in RS-232 serial port. (Assume no 1->3 adapter is used).

   mknod /dev/procomm c 1 0x090004

2. A terminal connected to port 3 of a 4-channel MUX at select code 29.

3. The "raw" I/O to the root disk. (Assume high-speed HP-IB.)

4. The "raw" and "block" I/O to cartridge tape at HP-IB address 3. (Assume standard-speed HP-IB.)

5. A SCSI MO disk at select code 14 with bus address of 2.

6. A CD-ROM drive at select code 8 with bus address of 4 using an HP-IB interface.

7. An 8-channel MUX at select codes 28 and 29 with a dial-out modem at port 0 and a dial-in modem at port 1. Assume the ports are in the second 4 port group.

8. Both right and left flexible disk drives (character *and* block) at select code 7, bus address 2.

9. The 7980 magnetic tape at select code 7, bus address 4, 6250 bpi, no rewind on close in AT&T compatibility mode.



9c

## CD-19. LAB: Device Files (Series 700)

### Directions

For this exercise it is *not* necessary to actually make the required device files, as there may be conflicts with existing files on your training system. Your instructor will indicate if any of the device files are to be created on your training systems.

Write out the command you would use to create device files for the following devices:

1. A direct terminal connected to the first built-in RS-232 serial port on the core I/O board.

2. The "raw" I/O to the root disk.

3. The "raw" I/O to DDS format DAT at bus address 4 on the core I/O board.

4. A MO disk in slot 1 on the EISA bus adapter.

5. A parallel port on the core I/O board.

6. The DDS format DAT using the core I/O board at bus address 4, no rewind on close in AT&T compatibility mode.

7. The DDS format DAT using slot 1 in the EISA bus adapter I/O board at bus address 4, no rewind on close in Berkeley compatibility mode.

8. A modem connected to the first built-in RS-232 serial port on the core I/O board.

# Module SS — System Startup

## Objectives

Upon completion of this module, you will be able to:

- Explain the purpose and syntax of the `/etc/inittab` file.
- Add an entry to `/etc/inittab`.
- Explain how `init` reads the `inittab` file.
- Identify the default run-levels supplied with an HP-UX system.
- Change the default system run-level.
- Explain the function of the `getty` process.
- Start a `getty` running on a terminal port.
- Describe the login process.
- Use the `tset` command to initialize terminal characteristics.



92

# Module SS — System Startup

## SS-1. SLIDE: Review of Boot ROM Startup

55

### Review of Boot ROM Startup

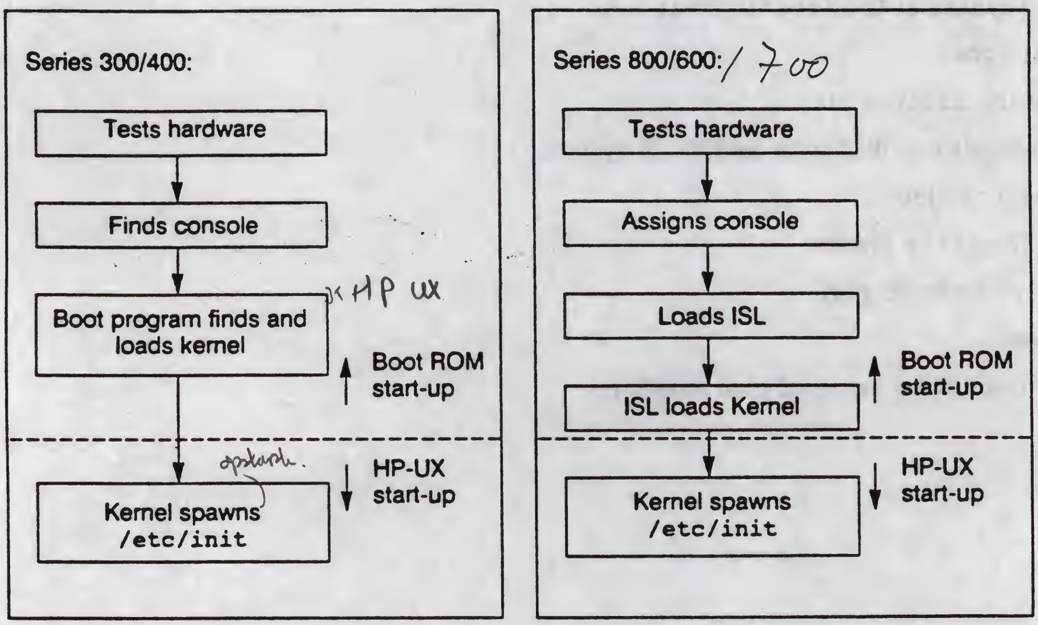


51436/51482 SS-1. 124 

## Student Notes

System startup occurs in 2 main phases:

1. the boot ROM startup sequence
2. the HP-UX startup sequence

We have looked at the boot ROM startup sequence which ends when the kernel takes control. Once the kernel takes control, the HP-UX startup sequence begins. This module is devoted to the HP-UX startup sequence.

## SS-2. SLIDE: HP-UX Startup Sequence and init

56

### HP-UX Startup Sequence and init

Once the HP-UX kernel takes control from the boot ROM, it starts the /etc/init process which:

- Spawns all other processes (PID = 1)
- Controls the run-level of the system
- Extracts arguments from /etc/inittab file

51436/51482 SS-2. 125 

## Student Notes

Once the HP-UX kernel takes control, it starts the /etc/init process. The init process has a process ID (PID) of 1 and has no parent. The init process reads the /etc/inittab configuration file, which defines the environment for the normal operating run-level (run-level 2 as shipped). The /etc/inittab file contains entries for all run-levels.



93

## SS-3. SLIDE: Run-Levels

57

### Run-Levels

- At all times, HP-UX is in a particular run-level
- Pre-defined run-levels are shipped with the system
- The run-level in which you system boots is defined by the `initdefault` entry in `/etc/inittab` (typically 2)

  `init:2:initdefault`
- The run-level can be changed with the `init` command
- Use the `who -r` command to determine the run-level

```
# who -r
.       system boot   Sep 22 12:52     2     0     S
```

Run level — teller — andre nivå

51436/51482 SS-3. 126 

## Student Notes

At all times, the HP-UX system is in a particular run-level. A **run-level** is a system state in which a specific set of processes is allowed to run. This set of processes is defined in the **/etc/inittab** file for each run-level. Run levels correspond to the following numbers/letter: 0, 1, 2, 3, 4, 5, 6, s, and S. You can define (or change) the run-levels 1-6. You cannot change run-levels 0, s, or S.

Your system comes with several pre-defined run-levels: run- level 2, run-level 0, run-level s, and run-level S.

Run-level 2 — is the multi-user run-level meaning that the system is available to multiple users. It is also the initdefault run-level (as shipped). **initdefault** is the run-level in which your system automatically boots.

Run-level 0 — is a special run-level reserved for system installation. You should not run in run-level 0.

Run-level s — is a special run-level reserved for system administration tasks. It is also referred to as single-user run-level meaning it is reserved for a single user, typically, the system

administrator. For example, shutting down the system (`/etc/shutdown`) brings you to run-level s.

Run-level S is similar to run-level s. With `init s` only the physical system console has access to the operating system, whereas `init S` (capital S) switches the capabilities of the system console to the terminal where you are logged in, thus making it the virtual system console.

The remaining run-levels can be designed by the system administrator. For example, the administrator may wish to define a run-level where only certain processes are allowed to run. We will talk more about designing your own run- level after we discuss `/etc/inittab`.

You can change the run-level of the system with the `init` command. Invoking `init` with an argument causes `init` to change the run-level of the system to the level specified by the argument. `init` scans `/etc/inittab` for all entries matching the new run-level (including those entries that are valid for all run-levels) and executes the commands associated with the entries. For example, the convention on most HP-UX systems is that run-level 2 is used for multi-user operation. So, to change the system from single-user to multi-user mode, the administrator would enter:

```
init 2
```

Whenever the run-level of the system is changed, any process which does not have an entry for the new run-level is sent a warning signal and then, after a 20 second grace period, is killed.

We have seen that the run-level of an HP-UX system is controlled by `init`. The actions of `init` are in turn controlled by a configuration file called `/etc/inittab`. Next, we will look at the contents and format of this file and how `init` utilizes this information to control the run-level of the machine.

Who ?

## SS-4. SLIDE: Fields in the inittab File

58

### Fields in the inittab File

Format:

```
id:rstate:action:process
```

Examples:

```
is:2:initdefault:

03:236:respawn:/etc/getty tty0p4 9600 #office L12
```

Baud Rate

51436/51482 SS-4. 127 

## Student Notes

Each of the fields in the `inittab` file is defined as follows:

*id* unique label which identifies entry (up to four characters).

*rstate* defines run-levels in which the entry will be processed.

*action* a keyword which defines how to execute the *process* (or program).

*process* shell command to be run if the entry's *rstate* matches the run-level and/or the *action* field indicates such action.

In the first example, the `initdefault` action causes the initial (default) run-level to be the value of the *rstate* field. The `initdefault` entry does not require a *process*. We looked at this entry briefly on the previous slide.

In the second example, each of the fields can be explained as follows:

| | |
|---|---|
| 03 | is the label, a unique number that identifies the entry in `/etc/inittab`. |
| 236 | are the run levels. When the current system run level matches the run level in an entry, the program associated with it is executed. |
| respawn | is the action that the system should take with the program. |
| /etc/getty tty0p4 9000 | is the program, the part of the entry that will be passed to the shell for execution. |
| #office L12 | is a comment. Comments must be preceded by a '#'. |

Each terminal or RS-232 port used as an incoming terminal device must have an **inittab** entry similar to the one above. We will discuss this in more detail in a short while.

---

**Note** You may have multiple entries in the *rstate* field. If the *rstate* field is empty, the entry is valid for all run-levels.

---



95

## SS-5. SLIDE: inittab Example— Creating New Run-Levels

59

### inittab Example

```
init:  4:initdefault:
              .
              .
              .
co:    :respawn:/etc/getty -h console console
01:234:respawn:/etc/getty tty0p1 9600#faculty 1
02:234:respawn:/etc/getty tty0p2 9600#faculty 2
05:23 :respawn:/etc/getty tty0p5 9600#grad 1
06:23 :respawn:/etc/getty tty0p6 9600#grad 2
08:2  :respawn:/etc/getty tty0p8 9600#student 1
09:2  :respawn:/etc/getty tty0p9 9600#student 2
```

51436/51482 SS-5 128 

## Student Notes

Assume there are terminals for three groups of people: undergraduate students, graduate students, and faculty. At all times, you want the faculty members to have access to the system. But there are times that you want to restrict undergraduate and/or graduate students from having access.

Case 1 — Undergraduate students are restricted from access. Both graduate and faculty have access.

Case 2 — Undergraduate and graduate students are restricted from access. Only the faculty has access to the system.

To do this, you would create 3 run-levels as shown on the slide. Run-level 2 allows everyone to access the system. Run-level 3 allows only graduate and faculty to access the system. Run-level 4 allows only the faculty members to access the system. The system comes up in run-level 4. To allow everyone access you would simply change the run-level of the system to 2.

If you wanted the system to come up so that, initially, only the faculty has access, you could change the `initdefault` entry:

```
init:4:initdefault
```



96

## SS-6. SLIDE: Typical `inittab` File

6 o

### Typical `inittab` File

```
init:2:initdefault:
brc1::bootwait:/etc/bcheckrc </dev/console >/dev/console 2>&1
brc2::bootwait:/etc/brc 1>/dev/console 2>&1
link::wait:/bin/sh -c "rm -f /dev/syscon;\
     /bin/ln /dev/systty /dev/syscon" >/dev/console 2>&1
cwrt::bootwait:/bin/cat /etc/copyright >/dev/syscon
rc::wait:/etc/rc </dev/console >/dev/console 2>&1
powf::powerwait:/etc/powerfail >/dev/console 2>&1
lp::off:nohup sleep 999999999 </dev/lp & stty 9600 < /dev/lp
cons::respawn:/etc/getty console console
01:23:respawn:/etc/getty -h tty0p1 9600# laboratory
02:2:respawn:/etc/getty -h tty0p2 9600# office K17
03:2:respawn:/etc/getty -h tty0p3 9600# office M17
04:2:respawn:/etc/getty -h tty0p4 9600# office L13
05:2:off:/etc/getty -h tty0p5 9600# printer port
pf::powerfail:/etc/powerfail 1>/dev/console 2>&1
```

51436/51482 SS-6. 129 

## Student Notes

These are entries in a typical **inittab** file. The file that comes with your system will be a little different. It is your responsibility as system administrator to modify this file as needed.

Let's look at the third and the fourth fields of the **inittab** file in more detail.

The third field, or the *action* field, contains a keyword which tells **init** how to execute the program specified in the fourth field. This allows you the ability to tailor the way in which processes are executed. For example, you can specify whether you want **init** to wait for the process to complete before starting another, or go ahead and start a new process while the first is still running.

There are many keywords allowable in the *action* field. The most important ones are listed below. For more information see *INITTAB(4)* in the *HP-UX Reference* manual.

initdefault — Causes the initial (default) run-level to be the value of the *rstate* field. If more than one run- level is specified in *rstate*, init uses the highest specified run-level.

| | |
|---|---|
| wait | On entering the run-level that matches the *rstate* field of this entry, run *process* and wait for it to die before reading the next entry. |
| boot | Run the command specified in the *process* field at boot-time only. Do not wait for *process* to die before reading the next entry. |
| bootwait | Run the command specified in the *process* field at boot-time only. Wait for *process* to die before reading the next entry. |
| respawn | On entering the run-level that matches the *rstate* field of this entry, run *process* if it is not already running. Do not wait for *process* to die before reading the next entry. If/when the process dies, run it again. |
| off | When the run-level is entered, if the *process* is running, `init` will send *process* a warning signal and then wait 20 seconds before killing it. If *process* is not running, the entry is ignored. Thus `off` is also used to deactivate an entry for some time. |

The fourth field of `inittab` is the *process* or *program* field. This is the shell command `init` will execute if the entry's *rstate* matches the run-level and/or the *action* field indicates such action. The command in this field is automatically `exec`'d and passed to a child shell as `sh -c 'exec` *command*'. A comment can be inserted in this field by preceding the comment with a "#".

For example, the `inittab` line

```
02:2:respawn:/etc/getty -h tty0p2 9600  # office K17
```

creates a `getty` child process, with `init` as the parent process.

## SS-7. SLIDE: init Actions After Booting

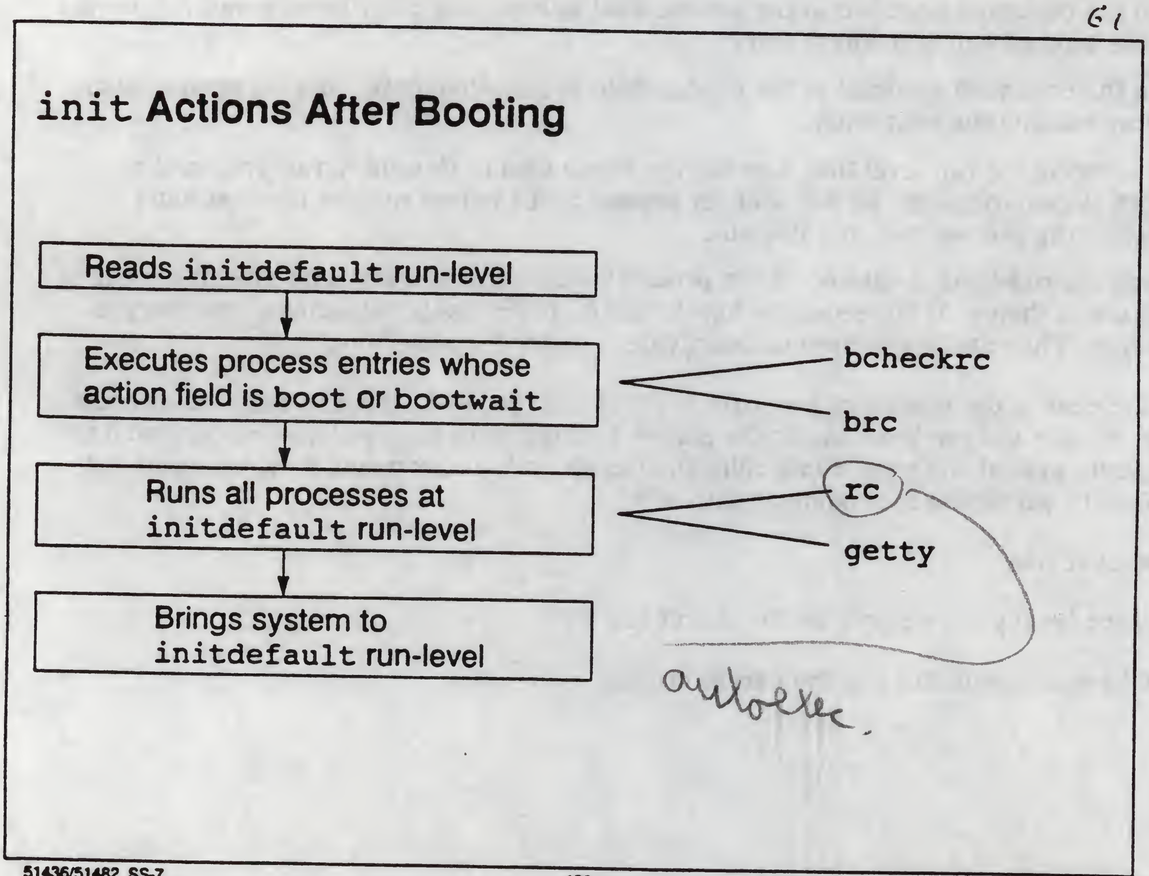


51436/51482 SS-7. 130 

## Student Notes

The first thing **init** does is scan **inittab** for an **initdefault** entry in the *action* field. The run-level associated with this entry is the initial run- level **init** will enter. If there is no entry for **initdefault**, **init** will prompt the administrator to specify a run-level to enter.

Next **init** will scan **inittab** for all entries marked **boot** or **bootwait** in the *action* field. Any commands associated with these entries are executed. In our sample **inittab** file, the programs **/etc/bcheckrc** and **/etc/brc** are run.

Next **init** calls a new program called **ioinit**. This program is new at 8.0 and it is part of the new auto-configurator. It scans the I/O backplane for new devices that need to be configured. After they have been automatically configured **init** continues on.

Following **ioinit init** executes all processes associated with the **initdefault** run-level. Typically, for a multiuser run-level (2), the logical (or virtual) system console is reset to the physical system console, the **/etc/rc** script is executed, and all **gettys** are started.

## SS-8. SLIDE:The /etc/bcheckrc Program

62

### The /etc/bcheckrc Program

```
brcl::bootwait:/etc/bcheckrc </dev/console >/dev/console 2>&1
```

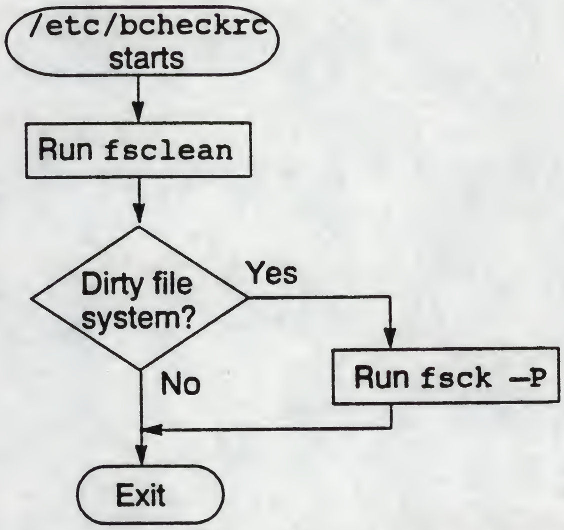


51436/51482 SS-7. 131 

## Student Notes

The **/etc/bcheckrc**(Boot CHECK Run Command) program checks to see if the system was properly shutdown. To determine if the system was properly shutdown, **bcheckrc** calls the **fsclean** program. **fsclean** checks each file system of type **hfs** in **/etc/checklist** to see if there might be a consistency problem. To do this, **fsclean** looks at a flag called the **clean byte** in the primary super block of each file system. When a file system is created, the clean byte flag is set to FS_CLEAN. When the file system is mounted (using the **mount** command), the clean byte flag is set to FS_OK. During normal shutdown (that is, during the execution of the **reboot** or **shutdown** command), the clean byte is reset to FS_CLEAN. So, under normal conditions, the file system can be unmounted and set to FS_CLEAN, or mounted and set to FS_OK.

If, when **fsclean** checks the clean byte, it finds the file system is unmounted and set to FS_OK, then the file system might be in an inconsistent state (due to a crash or other incorrect shutdown). In this case, **bcheckrc** will run **fsck** automatically using the **preen** mode. This will correct most errors found.



98

One note here, at 8.0 the root file system, or the primary file system is checked before `init` runs. So bcheckrc'' will check all file systems, except the root file system that has already been checked.

We will cover `reboot`,`shutdown`, `mount`, and `fsck` in later modules. If this does not sink in now, make a note of it and come back and reread this after we cover the module on Creating and Maintaining File Systems.

## SS-9. SLIDE:The /etc/brc Program

63

### The /etc/brc Program

`brc2::bootwait:/etc/brc >/dev/console 2>&1`

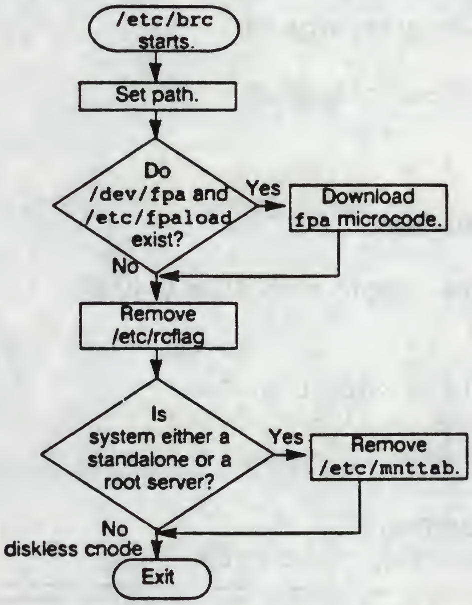


51436/51482 SS-7. 132 

## Student Notes

The **/etc/brc** program does the following tasks:

- Set the system PATH variable, which defines the default command search path for all users.
- Downloads the floating point microcode if your system is set up for the floating point accelerator.
- Removes the **/etc/rcflag** file (used later as a check for startup condition).
- Checks for, and removes, the **/etc/mnttab** file on a standalone or cluster server system. The **/etc/mnttab** file contains a list of mounted file systems. This file will be re-created later when **/etc/rc** re-mounts the file system.



99

## SS-10. SLIDE: The /etc/rc Program

69

PS — alle & taken

### The /etc/rc rc Program

rc::wait:/etc/rc </dev/console >/dev/console 2>&1

- Executes when new run-level is invoked
- Exact tasks depend on status of system
- For all run-levels:
  - Sets TZ (time zone) variable
- When moving from single-user to multi-user (s to 2):
  - Mounts file systems
  - Activates swapping to all swap devices
  - Starts syncer → schijfbuffer
  - Starts cron → taken
  - Starts lpsched → printer regelaar
  - Starts networking daemons
  - Performs other "housekeeping" chores

51436/51482 SS-7. 133 

## Student Notes

As we have seen, whenever the run-level of the machine is changed with init the file /etc/inittab is read for entries matching the new run-level. One of the entries in inittab may be to invoke /etc/rc. As shipped, this file is invoked every time the run-level of the HP-UX system is changed.

The /etc/rc script consists of a main script program and several shell functions (a shell function is essentially a subroutine in a shell program). The main program checks to see if it is boot time (known by the presence or absence of /etc/rcflag). If so, it determines if you are running as a standalone system, a cluster server, or a cluster client. It then calls shell functions applicable to your system. When finished, it calls a shell function called localrc. This is the shell function you should use to customize the /etc/rc script. It should contain any tasks you wish to perform that are not part of the standard /etc/rc functions. In the localrc shell function, you can add commands you wish to perform every time the system is booted or whenever there is a change in run-level which init does not handle.

# Module SS — System Startup

The table belows shows what functions `/etc/rc` performs for each system state. Note that some commands may not be available or installed. `/etc/rc` checks for the existence of all commands before attempting to run them.

**Table SS-1.**

| Function | Standalone | Root Server | Diskless Cnode |
|---|---|---|---|
| set host name | X | X | X |
| initialization: set TZ and other variables | X | X | X |
| local functions: any functions you put into the script | X | X | X |
| set the date | X | X | |
| create the `/etc/setmnt` file | X | X | |
| mount all `hfs` volumes listed in `/etc/checklist` | X | X | |
| start the `syncer` | X | X | |
| start the `lp` scheduler | X | X | |
| clean uucp and editor files | X | X | |
| start networking | X | X | X |
| start swapping to all swap devices in `/etc/checklist` | X | X | X |
| start `cron` | X | X | X |
| start pty allocation daemon | X | X | X |
| start `vtdaemon` | X | X | X |
| list files found in `/tmp` and `/usr/tmp` | X | X | X |
| clean logging files | X | X | X |
| start CSPs | | X | X |
| start the remote boot daemon | | X | |
| start auditing processes | X | X | |
| save a core image of a previously crashed system | X | X | X |
| start the diagnostic logging for I/O subsystem (800 only) | X | X | X |
| start the logging system messages (800 only) | X | X | X |

HP-UX

init & ← inittab | Check's /etc/RC

## SS-11. SLIDE: The /etc/getty Command

65

### The getty Process

- Normally invoked by init via the /etc/inittab file
- First command executed for each login
- Operates as follows:

1. Displays the contents of /etc/issue
2. Issues login: prompt
3. Waits for you to type something then reads login name
4. Establishes speed and case
5. Invokes /bin/login passing it the login name you typed

/etc/getty [ -h ] [ -t *timeout* ] *line* [*speed* ]

-t *timeout* is the timeout value (seconds)
-h forces line hangup after timeout
*line* is the tty line in /dev
*speed* is the label to speed definition in /etc/gettydefs

51436/51482 SS-8. 134 

## Student Notes

On most HP-UX systems, run-level 2 signifies multi-user operation. In a multi-user run-level, a getty process is running on each port in which a user can log in. The getty process is normally invoked by init according to an entry in inittab.

The purpose of getty is to set terminal options, print the contents of the /etc/issue file (if it exists), print a login prompt, wait for input to that prompt, and, following a response by a user, exec /bin/login.

You must supply getty with the device file name of the line on which it is to run. This device file should exist in the /dev directory.

Frequently, getty is invoked with a -t option followed by an integer value representing seconds. If this option is specified, getty opens a line and if nothing is typed in the number of seconds specified, getty exits.

getty may also be invoked with a -h option. This option is used primarily with modem connection lines. If specified, getty drops carrier when a user logs off. Another login prompt is not issued and the user must dial in again to establish a new connection.

getty also has a speed option. If specified, it serves as a label in to the /etc/gettydefs file. The definition in /etc/gettydefs instructs getty at what speed to run, what to use as a login prompt, what to set as initial tty line settings, and at what speed to try next if the initial speed is inappropriate. With HP-UX, a speed entry of 9600 should be used for terminals connected directly to the system. For dial-up ports, a label corresponding to the appropriate speed for the dial-up line should be used. If a speed value is not designated, a default of 300 baud is used.

The *action* field for a getty entry in inittab is usually respawn. Thus, whenever the getty process terminates, usually when the user logs out, a "wake-up" signal is sent to init. init immediately forks a new getty process. The result is that another login: prompt appears on the terminal connected to that port.

The actions of getty are extensive and complex. For more information, see *GETTY(1M)* in the *HP-UX Reference* manual.

**Note** As shipped, /etc/inittab invokes /etc/getty only for the system console in run-level 2. If your system has additional terminals on which you wish to support logins, you must add the appropriate getty entries to /etc/inittab. (SAM automatically creates these entries when you use it to add terminals.)

## SS-12. SLIDE: The gettydefs File

66

### The /etc/gettydefs File

**label# initial-flags # final-flags #login-prompt# ext-label**

- label
  - Identifies the entry
  - Matches against getty speed argument in /etc/inittab
- initial-flags
  - Initial line and terminal settings
  - Speed must be specified
- final-flags
  - Final line and terminal settings
  - Speed must be specified
- login-prompt
  - Initial login-prompt printed on the terminal
- Next-label
  - Entry to try next if "break" is typed

51436/51482 SS-9. 135 

## Student Notes

The file **/etc/gettydefs** contains information used by **getty** to set up speed and terminal settings for a line. Each entry in the **gettydefs** file contains a series of hash mark ("#") separated fields having the following format:

**label# initial_flags # final_flags #login_prompt#next_label**

The meaning of each field is as follows:

label — The string against which **getty** tries to match its second argument. It is often the speed, such as 1200, at which the terminal is supposed to run, but it need not be.

initial-flags — These flags are the initial **ioctl(2)** settings to which the terminal is to be set if a terminal type is not specified to **getty**. The flags that **getty** understands are the same as the ones listed in **/usr/include/sys/termio.h**. Normally only the speed flag is required in the *initial-flags*. **getty** automatically sets the terminal to raw input mode

and takes care of most of the other flags. The *initial-flags* settings remain in effect until `getty` executes `/bin/login`.

final-flags — These flags take the same values as the *initial-flags* and are set just prior to when `getty` executes `login`. The speed flag is required. The composite flag SANE takes care of most of the other flags that need to be set so that the processor and terminal are communicating in a rational fashion. The other two commonly specified *final-flags* are TAB3, so that tabs are sent to the terminal as spaces, and HUPCL, so that the line is hung up on the final close.

login-prompt — This entire field is printed as the *login-prompt*. Unlike the above fields where white space is ignored, white space is included in this field.

next-label — If this entry does not specify the desired speed, indicated by the users typing a (break) character, then `getty` will search for the entry with *next-label* as its *label* field and set up the terminal for those settings. Usually, a series of speeds are linked together into a closed set. For example, 2400 is linked to 1200, which in turn is linked to 300, which finally is linked to 2400.

The syntax of the `gettydefs` file must be exact. Spaces within each entry must appear as shown on the slide and a blank line must follow each entry (apart from the very last line).

Two sample entries in the file `/etc/gettydefs` might look like these:

```
9600    # B9600 HPUCL SANE CS7 PARENB ISTRIP IXANY TAB3
        # B9600       SANE CS7 PARENB ISTRIP IXANY TAB3
        #login: #300

console# B9600 SANE CLOCAL CS7 PARENB ISTRIP IXANY TAB3 HUPCL
        # B9600 SANE CLOCAL CS7 PARENB ISTRIP IXANY TAB3 UUPCL
        # Console login: #console
```

For more information on the `gettydefs` file see *GETTYDEFS(4)*, *STTY(1)* and *TERMIO(7)* in the *HP-UX Reference* manual.

**Note** — It is strongly recommended that after making or modifying `/etc/gettydefs`, it be run through `getty` with the `-c` option (for check) to be sure there are no errors.

The following procedure is recommended:

```
# cd /etc
# cp gettydefs gettydefs.new
# vi gettydefs.new
    .
    .
    .
# getty -c /etc/gettydefs.new | more
# mv gettydefs.new gettydefs
```

## SS-13. SLIDE: Adding a Terminal With SAM

67

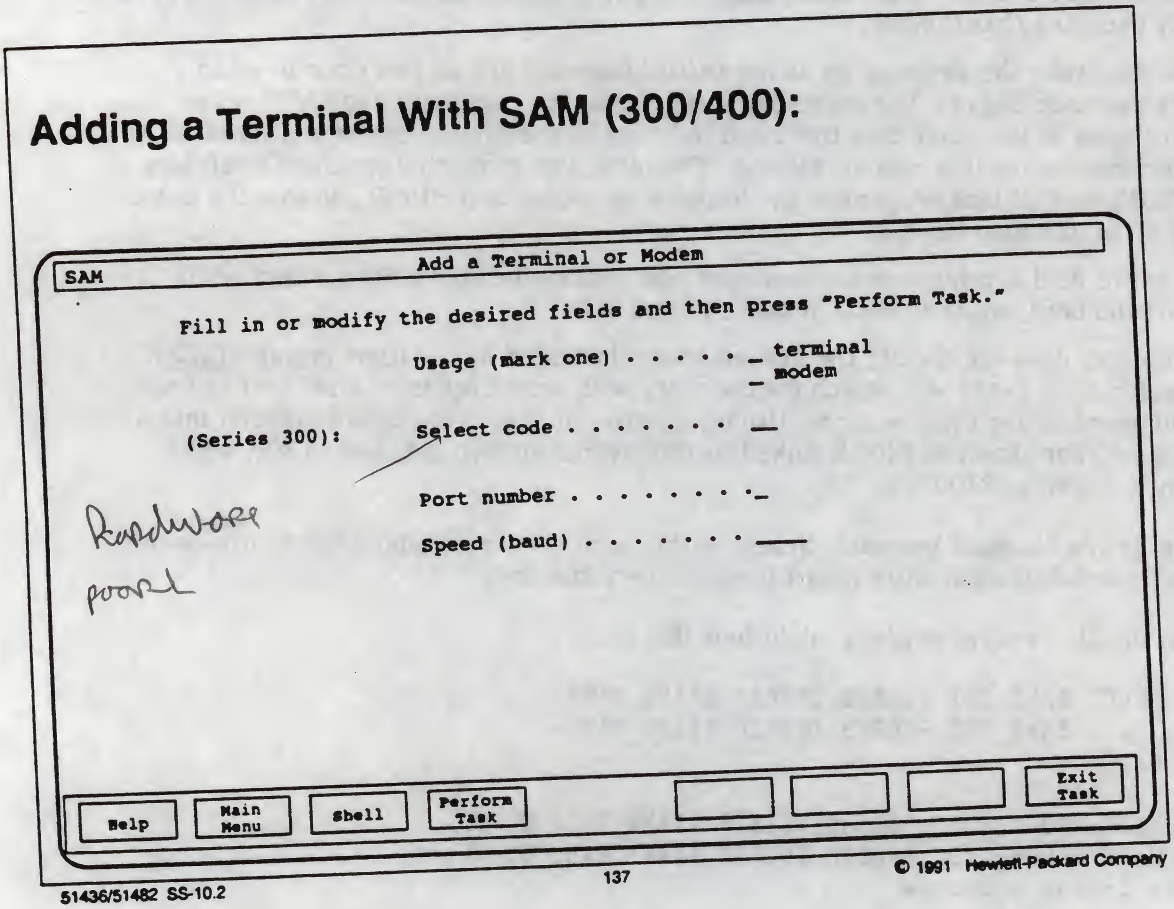


## Student Notes

Before you use SAM to add your terminal, you must connect the terminal to your system. Make a note of the select code (Series 300) or logical unit number (Series 800) and port number of your new terminal. Once the terminal is physically connected to the system, you can use SAM to make HP-UX recognize the terminal.

To add a terminal with SAM, select "Peripheral Devices ->" from the Main Menu, then select "Add a Terminal or Modem . . ." That will bring you to the menu shown on the slide. Fill in the fields with the appropriate information and then press (Perform Task). SAM will give you this message:

```
Adding Terminal
```

Then you will see:

```
Device Added
```

```
-- Press the space bar to continue. --
```

Press the space bar, and you will get the message:

```
Would you like to add another device? (y or n)
```

If you are finished, answer `n`. SAM created the necessary devices files and added the appropriate entries to the **/etc/inittab** file for you.

### To Add a Terminal Manually

If you want to add the terminal manually, you would need to perform the following steps:

1. Connect the terminal.
2. Make the appropriate device file. (Creating device files is covered in a later module.) Series 300 example:

   ```
   # mknod /dev/tty13 c 1 0x0d0003
   ```

   Series 800 example:

   ```
   # mksf -d mux0 -l 1 -p 3 -h  /dev/tty13
   ```

3. Check **/etc/gettydefs** and add baud rate chain for new terminal if necessary, or use existing baud rate chain.
4. Edit **/etc/inittab** and add a **getty** for new terminal.

   ```
   13:2:respawn:/etc/getty tty02 9600
   ```

5. Invoke `init q`.
6. Invoke **ps -ef** to see if the **getty** is running.



103

## SS-14. SLIDE: Terminal Setup Characteristics

6 8

stty
Kap

### Terminal Setup Characteristics

- TERM environment variable contains the terminal type. For example: file terminfo = bin file

  ```
  TERM = (hp)
  ```

- tset, sets the value of TERM, initializes terminal characteristics, and configures the terminal port. Examples:

  ```
  $ tset
  Erase is Backspace
  Kill is Ctrl-U

  $ eval ' tset -s -Q -m ':300h' '
  ```

- /etc/ttytype is a database of terminal types by port. For example:

  ```
  300h console
  70092 tty0p1
  dialin ttyd00
  ```

- stty allows you to modify and display terminal characteristics

51436/51482 SS-11. 138 

## Student Notes

To communicate effectively with your terminal, HP-UX must know the type of terminal or graphics display you are using. The TERM environment variable supplies the terminal type to HP-UX and can be set with the `tset` command.

The default local login script prompts you to enter your terminal type as follows:

```
TERM = (hp)
```

Pressing (Return), sets the TERM environment variable to hp, the default value. This value works with HP terminals, but it may not let you take full advantage of your terminal or graphics display features. Entering a different value sets the TERM environment variable to that value.

HP-UX supports many terminal types. The `/usr/lib/terminfo` database tells HP-UX how to communicate with each terminal type. When you assign a value to TERM, the value must equal a value in the `terminfo` database. See the next slide.

The `tset` command allows you to set the value of TERM and initialize your terminal characteristics. Some common user errors, such as trying to display the contents of a binary file (for example, an executable file) on a terminal, might leave the terminal in an unusable state. You can use `tset` to reset your terminal (shown on slide).

```
tset
```

`tset` simply reads the terminal type out of the environment variable TERM and reinitializes the terminal.

If you always log in using the same terminal type, you could change your local login script to eliminate the TERM prompt. In the local script, this command displays the TERM prompt:

```
eval ` tset -s -Q -m ':?hp' `
```

To customize the above command, replace ?hp with the value of TERM. For example, the following command (shown on slide also) initializes your terminal as a high-resolution graphics display (300h), but the TERM prompt itself does not display:

```
eval ` tset -s -Q -m ':300h' `
```

If you use more than one type of terminal (such as one at work and one at home), you could modify your `tset` command to include multiple terminal types.

The `stty` command can be used to change various terminal characteristics, such as baud rate, erase and kill characters, etc. For example, the following stty command resets the erase and kill characters to # and @, respectively:

```
stty erase # kill @
```

`stty` also displays terminal settings if invoked with the -a option.

`ttytype` is a database containing, for each tty port on the system, the kind of terminal that is attached to that port. There is one line per port. It contains the terminal kind, a space, and the name of the tty. This information is read by `tset` and `login` to initialize TERM at login time. `/etc/ttytype` should be edited by the system administrator to match the actual port configuration.

## SS-15. SLIDE: `terminfo` Data Base

69

**`terminfo` Data Base**

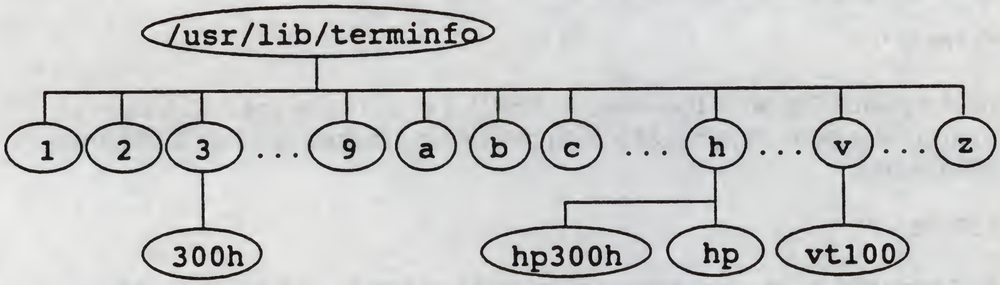


51436/51482 SS-12. 139 

## Student Notes

The contents of the environment variable TERM is a pointer to a file in the terminal capability database `terminfo`. Terminals are described in `terminfo` by giving a set of capabilities which they have and by describing how operations are performed.

For the sake of a faster access to the `terminfo` files, an additional level of directories was added. These directories are named after the initial letters of the file names.

The `terminfo` database contains files describing HP terminals and many commonly used non-HP terminal. The set of files is updated by HP. Usually a system administrator need not modify these files. If you find you need to modify a terminfo file, translate it from an internal, compiled format into an open, source format with the `untic` command. The `tic` command compiles the modified source file and places the compiled file back into the `terminfo` database.

## SS-16. SLIDE: What Happens At Login

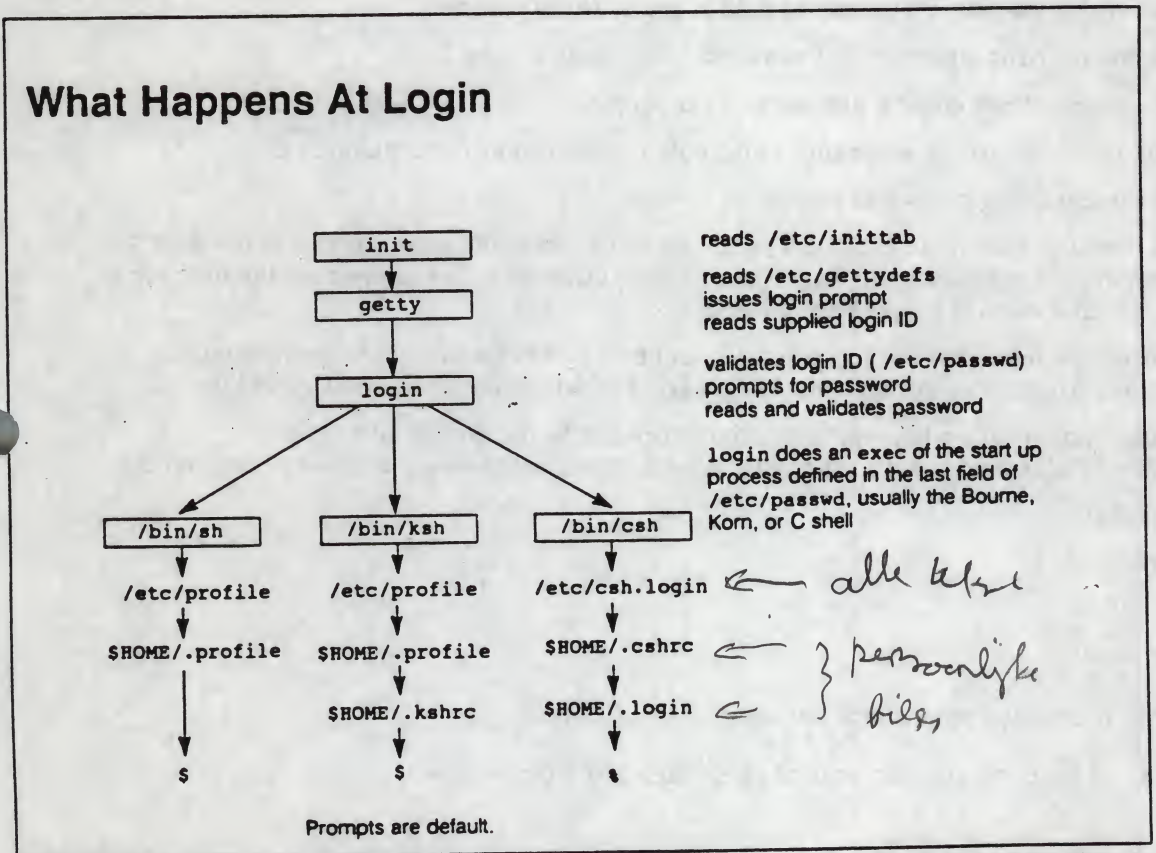


51436/51482 SS-13. 140 © 1991 Hewlett-Packard Company

### Student Notes

Allowing users to log into an HP-UX system is a three step procedure. The steps are:

`/etc/getty -> /bin/login -> /bin/ksh, keysh, csh or sh`

We have already discussed the operations of `getty`. Now we will discuss `login`. The following steps describe what `login` does.

1. `login` searches the `/etc/passwd` file for the username.
   a. if the username exits, `login` goes to step 2.
   b. if the username does not exist, `login`:
      i. Prompts the user for a password (`Password:`). This makes it difficult for an intruder to find and use a valid username.
      ii. Displays the message `Invalid login`.

   iii. Updates the `/etc/btmp` file (if it exists). `/etc/btmp` keeps track of invalid login attempts.

   iv. If this is the user's third consecutive invalid login attempt, `login` exits; otherwise, `login` prompts the user (`login:`) for a username and repeats step 1.

2. `login` checks to see if the username's password field is set in `/etc/passwd`.
   a. If so, it prompts the user for a password (`Password:`) and goes to step 3.
   b. If not, then the user need not enter a password; go to step 4.
3. `login` compares the password to the username's encrypted password in `/etc/password`.
   a. If the password matches, `login` goes to step 4.
   b. If the password does not match, `login` displays the message `Invalid login.` If this is the user's third consecutive invalid login attempt, `login` terminates; otherwise, `login` prompts the user for a username (`login:`) and control passes back to step 1.
4. `login` sets the username's numeric user ID, group ID, and home directory from the corresponding fields in `/etc/passwd`. `login` also updates the `/etc/wtmp` file, which keeps track of valid logins.
5. `login` runs (via `exec` system call) whatever command is present in the command field of `/etc/passwd`. Typically, this field is set to the path name of the shell the user wishes to use; that is:
   a. `/bin/ksh` (Korn shell)
   b. `/bin/sh` (Bourne shell)
   c. `/bin/csh` (C shell)
   d. `/bin/pam` (PAM shell)

   If the command field is empty, `login` starts a Bourne shell by default.

We'll look further into the function and contents of these files in a later module.

## SS-17. WORKSESSION: Review Questions

1. How is the run-level of the system controlled?

init ↵ inittab

who -r

2. What is the purpose of /etc/inittab?

door init gelezen
RC
getty

sulk = bRC

3. What is the difference between init s and init S?

4. Describe the entire process from boot all the way until the user gets his $ prompt.

5. Describe what the /etc/rc file does.

· netwerk deamons

6. What does getty do?

Wacht op login

gettydefs — term



1c6

7. How can `gettydefs` entries be set up to handle I/O from different speed terminals?

8. What does `tset` do?

tekor karakteristiko
ni terminfo

## SS-18. LAB: Lab Exercises (Series 300 Only)

1. Determine the default run-level for the system you are using.

2. Determine the current run-level for the system being used. How many times has the current state been entered previously? What was the previous state of your system?

3. Determine your terminal port.

4. Determine the `inittab` entry related to your terminal port.

5. Execute **`ksh`** to create a child shell process we can play with. Now execute **`unset TERM;vi dummy`**. What happens? Why?

6. Determine what processes are currently running on your system. Change the run-level from "2" (multi-user) to "1". What processes were killed?

7. Set up your system so that you have a getty running on a terminal port. Make the getty work so that someone may log in at baud rates of 300, 1200, 2400, or 9600. The person logging in should merely have to press break to get to the next sequence. Make the login prompt change depending on what baud rate is established. If a terminal is available, connect the device and check to see that what you did worked.

## SS-19. LAB: System States (series 800 Only)

1. Determine the default run-level for the system you are using.

2. Determine the current run-level for the system being used. How many times has the current state been entered previously? What was the previous state of your system?

3. Determine your terminal port.

4. Determine the `inittab` entry related to your terminal port.

Execute `ksh` to create a child shell process we can play with. Now execute `unset TERM;vi dummy`. What happens? Why?

6. Add an entry into the `/etc/rc` script. The entry should be something like "echo RON WAS HERE". Then watch the next time that the system is booted, your message should appear during the boot process.

# Module SD — System Shutdown

## Objectives

Upon completion of this module, you will be able to:

- Understand the importance of a proper system shutdown.
- Explain the differences between `shutdown` and `reboot`.
- Use the `shutdown` command.
- Use the `reboot` command.



109

## SD-1. SLIDE: Why Shutdown the System

71

### Why Shut Down the System

- To conduct administrative activities without user interference, such as:
  - Checking file systems
  - Unmounting file systems
  - Backing up file systems
  - Reconfiguring the kernel
- To halt the system, so it can be turned off
- The `shutdown` command is most commonly used to shutdown the computer system.

51436/51482 SD-1. 141 

## Student Notes

As we have seen, most of the time the HP-UX system will be in multi-user mode, allowing many user and system processes to run. There are occasions, however, when the administrator must change the run-level of the system from multi-user to single-user. For example, if the administrator wants to backup a file system, the users should not continue to work and possibly change files while the backup is occurring. The administrator should bring the system to single-user state before backing up. Once in single user mode, the administrator might want to halt the system completely or reboot the system.

**Halting** brings the system to a complete stop; in this state, the only way to restart the system is to cycle the power or reset the hardware. **Rebooting** brings the system to a complete stop, but then restarts the system as if you had booted it.

Whether to halt or reboot your system depends on why you want to shutdown in the first place. If you want to shut the computer off for an extended period of time (for example, to add new hardware or to leave the system off for a long weekend), then halting is appropriate. If you want to shutdown only to boot the system again (for example, to use a newly configured kernel), then rebooting is appropriate.

# Module SD — System Shutdown

There are several different ways to shutdown an HP-UX system. The administrator is responsible for making sure the system processes are halted in a consistent and orderly manner. *Pushing the power button is not the proper way to shutdown your system.* Typically, the `shutdown` command is used to bring an HP-UX system from multi-user mode to single-user mode.

---

**Note** If an HP-UX system is shutdown improperly, you run the risk of corrupting the file system.

---



) C

## SD-2. SLIDE: The shutdown Command

72

### The shutdown Command

Used to bring system from multi-user to single-user mode, and then, possibly, halt or reboot

- shutdown first checks to see if this user has permission to use the shutdown command. If the user does have permission, it then:
- Changes the current working directory to /
- Updates all Superblocks
- Warns the users currently logged into the system of the impending shutdown
- Custom user scripts in the /etc/shutdown.d directory are executed
- All daemons processes are halted
- All processes are killed
- Writes the contents of the I/O buffers to disk
- Unmounts file systems
- Puts the system in single user mode
- Calls /etc/reboot if applicable

Format:

/etc/shutdown [ -r | -h ] [ *grace* ]

- -r automatically reboots after reaching run-level s
- -h halts the system completely
- *grace* means wait *grace* seconds before shutting down

51436/51482 SD-2. 142 

## Student Notes

**shutdown** stops system activities in an orderly and consistent manner. After system activities are stopped, administrative activities can be performed, or the machine can be rebooted or turned off.

The format of the command is shown on the slide. If you invoke the command with no option, it brings the system into single-user mode. The -h option halts the system after it reaches single-user mode. The -r option reboots the system after it reaches single-user mode. The argument *grace* allows the administrator to specify a grace period. The value given is the number of seconds that shutdown will wait before terminating all the processes running on the system. The default *grace* value is 60 seconds.

To execute shutdown:

1. Execute the shutdown command. (Good practice dictates that this command be executed from the system console.)

```
# shutdown
```

If there are other users on the system, shutdown prompts to see whether you wish to send the standard broadcast message or enter your own message. If you elect to send your own message, type the message on the terminal when prompted, press (return) and then Ctrl-d to signify the end of the message. (If there are no users, shutdown will not broadcast a message.)

After the message, shutdown:

- Waits the specified (or default) number of seconds.
- Stops system programs such as accounting.
- Terminates all other processes (such as error reporting, line printer spooling, gettys and shells) running on the system.
- Changes the run-level to single-user.
- Executes sync to flush the system buffers.
- Unmounts any file systems other than root.

If shutdown was invoked with no options, the system will be brought down to a single-user state and the administrator can now do whatever activities are required. If an option was used with shutdown, the system will either reboot (-r) or halt (-h).

**Note** The /etc/shutdown file is NOT a shell script and CANNOT be modified by the system administrator.

Instead at 8.0 customized scripts are made available for the system administrator to use. There is a new directory called /etc/shutdown.d, and in this directory the system administrator is free to put as many scripts as he needs. The scripts are executed in (ASCII) order. All dot files and subdirectories are ignored. These scripts may be application shutdown scripts, such as database shutdown scripts, etc.

Also new at 8.0 is the /etc/shutdown.allow file. This file allows the system administrator to specify who can use the shutdown command. This means that the system administrator can allow other users, such as system operators, to be able to use shutdown.

Examples of shutdown:

This example will shutdown and then halt your system with no grace period:

```
$ shutdown -h 0
Do you want to send your own message? (y or n):   n
```

This example will shutdown and then reboot your system giving the users 5 minutes grace:

```
$ shutdown -r 300
Do you want to send your own message? (y or n):   y
Type your message followed by ctrl d....

Please log off now.  The system will be rebooted in 5minutes ...
```

```
Ctrl-d
```

(5 minutes later)

```
Do you want to continue? (y or n):   y
```

(Note: Not all messages generated by shutdown are shown. )

## SD-3. SLIDE: The `reboot` Command

73

### The `reboot` Command

Usually used to reboot or halt the system once in single-user mode

`/etc/reboot [-h | -r] [-n | -s] [-m` *mesg*`] [-t` *time*`] [-f` *lif*`]`

| | |
|---|---|
| -h | halt the system |
| -r | reboot the system automatically (default) |
| -n | no `sync` before halt or reboot |
| -s | `sync` before halt or reboot (default) |
| -m *mesg* | supply message to users |
| -t *time* | halt at specified *time* |
| -f *lif* | reboot from specified lif file |

51436/51482 SD-3. 143 

## Student Notes

If you used shutdown to bring the system down with no options, then the system remains in single-user run-level allowing you to perform various tasks. Once you have performed these tasks, you might want to reboot or halt the system. The best way to do this is with the `reboot` command.

The default action of the `reboot` command is to sync the disks and reboot the system. reboot should always be invoked with the system in a single-user run-level. The format of the command is shown on the slide. The options are explained in more detail below.

| | |
|---|---|
| `-h` | will shutdown the system and halt the CPU. |
| `-r` | option will reboot the system (default). |
| `-n` | prevents computer sync from executing before the system is rebooted or halted. |
| `-s` | invokes `sync` before rebooting or halting the system (default). |



112

**-t** *time*\ **allows you to specify what time to bring the system down.** *time* **can be the word** *now* **(indicating immediate shutdown) or a future time in one of two formats:** *+number* **and** *hour:min.* **\-m** *mesg*\ **displays** *message* **at the terminals of all users on the system at decreasing intervals as reboot** *time* **approaches. \-f** *lif*

reboots from the specified file. If the filename is an empty string (`-f ""`), the power-up search sequence is made for a system (allowing you to boot in attended mode if you wish).

## SD-4. SLIDE: shutdown and reboot Examples

74

### shutdown and reboot Examples

Shutdown, perform administrative tasks, and reboot system:

```
# shutdown
(perform tasks)
# reboot
```

To activate a newly configured kernel, shutdown with no grace period and automatically reboot:

```
# shutdown -r 0
```

To install an interface card, halt system giving users 5 minutes to log off:

```
# shutdown -h 300
```

51436/51482 SD-4. 144 

### Student Notes

So, how do you decide whether to use `shutdown` or `reboot`. We have given you some examples on the slide. The command you use generally depends on:

- Whether users are logged in
- How quickly you need to shutdown the system

The `shutdown` command shuts down more slowly than `reboot`, but more gracefully, It uses `kill -14` to kill running processes, which lets processes terminate naturally within a grace period. This is the safest way to shutdown and it helps ensure file system integrity.

It also displays messages directing users to log off within a specified grace period. You can specify the grace period when you invoke `shutdown`. It is typically used when:

- The system is in a multi-user state.
- The system administrator is not the only person logged in and using the system.

The `reboot` command normally shuts down all processes very quickly. It uses `kill -9` to kill any running processes. This can be dangerous (for example, cause loss of data) because `reboot` will shutdown processes immediately without letting them terminate normally. It is typically used when:

- The system is in run-level s.
- You need to bring the system down very quickly.

# Module SD — System Shutdown

## SD-5. WORKSESSION: Review Questions

1. While in multi-user mode, why should the `shutdown -h` command be used, rather than `reboot -h` for halting the system?

2. What does `reboot -n` do, and why would you want to do this?

3. Once the system is in a s (or S) run-level (after execution of `shutdown`) , what is the difference between `init 2` and `reboot`? We assume a `initdefault` entry of 2 in `/etc/inittab`.

4. If the System Administrator wants to customize the shutdown process, what should he do?

5. Is the System administrator the only person who can shutdown the system? If not how can he give this capability to another user?



114

## SD-6. LAB: Lab Exercises (Series 300 Only)

1. Shutdown your system immediately (0 seconds) to single-user mode using the `shutdown` command. Look at what processes are still running. Then reboot your system with the `reboot` command.

2. After rebooting in exercise #1, use `init s` to shutdown your system. Look at what processes are still running. How does this compare to the processes that were running after using the shutdown command in exercise #1? Reboot your system. (`reboot` or `shutdown -r 0` will work.)

3. Customize the shutdown procedure by adding your own scripts to the `/etc/shutdown.d` directory.

4. Set it up so that your User logon name can execute the `shutdown` command.

## SD-7. LAB: Lab Exercises (Series 800 Only)

1. Using the simulations made available by your instructor, run the simulation of system shutdown to see what happens when the system is shutdown.

2. Check to see what processes are currently running? Change the system run-level from multi-user to 1 with the `init 1` command. What processes are running? Reboot the system.

3. Now shut the system down to a single-user run-level with the `shutdown` command. What processes are still running?

4. Now change to run-level 2 with the `init 2` command. What processes are running? Finally, return to regular multi user operation.

5. Customize the shutdown procedure by adding your own scripts to the `/etc/shutdown.d` directory.

6. Set it up so that your User logon name can execute the `shutdown` command.

# Module FS — File System Structure

## Objectives

Upon completion of this module, you will be able to:

- List the three components of an HFS file system.
- List the dynamic information in a superblock.
- Describe a track, cylinder, and cylinder group.
- List the contents of a cylinder group.
- Describe the contents of an inode table.
- List the pieces of an inode.
- Illustrate the structure of an inode.
- Explain three different ways that an inode references a data block.
- Explain fragment and block allocation.
- Describe disk sectioning (Series 600/800 users only).



116

# Module FS — File System Structure

## FS-1. SLIDE: The HP-UX File System

75.

### The HP-UX File System

Based on 4.2 BSD UNIX file system also known as:

- High-Performance File System (hfs)
- Berkeley File System
- McKusick File System

Redesign of traditional AT&T filesystem to accomodate larger disks

Same user interface as traditional AT&T filesystem

51436/51482 FS-1

1



## Student Notes

The HP-UX file system is based on work done by Kirk McKusick at University of California at Berkeley. The traditional UNIX file system developed by AT&T was developed when disks were small. As disks sizes grew by orders of magnitude, the initial design of the traditional UNIX began to show its age in respect to performance.

McKusick solved this performance problem by redesigning the underlying data structures that support the file system. With this approach, the use interface was unchanged and as such no code conversion was necessary for applications developed on the traditional UNIX file system. The simple and elegant user interface of the traditional UNIX file system coupled with the performance increases of the Berkeley group make the hfs file system a logical choice for HP-UX. HP's version of the file system has been tuned further for HP hardware.

## FS-2. SLIDE: Disk Layout (Series 300/400)

76

### Disk Layout (Series 300/400)

Disks can hold:

- Boot area (8KB)
- File system
- Swap

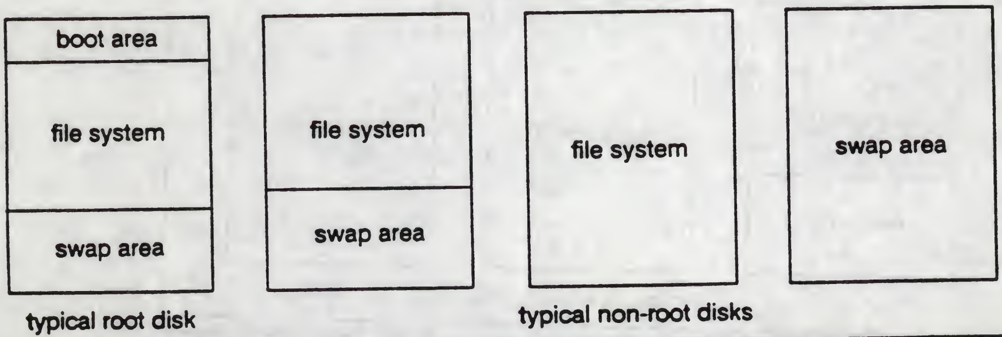


51436-51482 FS-2. 2 

## Student Notes

On a Series 300/400 HFS implementation, a hard disk may be comprised of a boot area, a file system area, and a swap area. The **root disk** will hold all three of these areas. (The root disk or system disk is the disk that holds the root file system.) At the start of the root disk is a 8KB area reserved for the boot. This 8KB area is a file system in its own right, albeit not a a UNIX file system. Its format is Logical Interchange Format or LIF. LIF is a very simple file system with simple data structures that can easily fit in a small area such as 8KB. Its simplicity makes it a good choice at boot up in order to keep the code size small in the boot ROM. This LIF volume actually holds, the secondary loader, the boot up code that will load in the HP-UX kernel, `/hp-ux`. This area is followed by the file system area for the root file system. The primary swap area is located just after the root file system.

Non-root disks will typically contain a single swap area, a single file system, or a combination of a both a single swap area and a single file system area. On the Series 300/400 computer there is only one file system area per physical disk.



117

## FS-3. SLIDE: Disk Layout (Series 700)

77

### Disk Layout (Series 700)

Disks can hold:

- LIF directory (8KB)
- File system
- Swap
- Boot area (2MB)

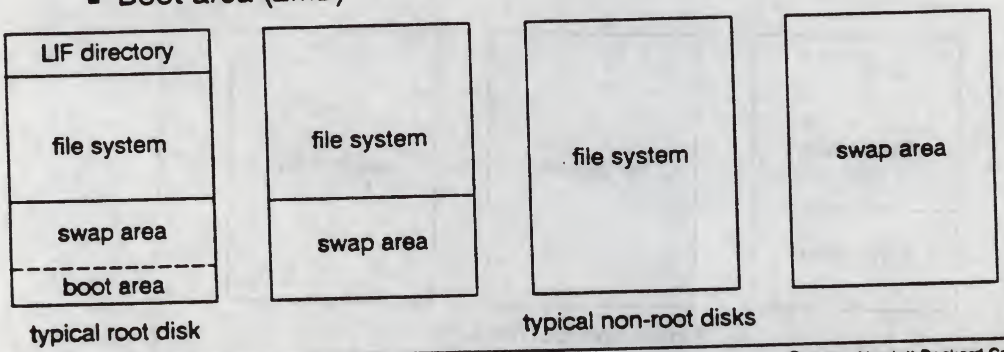


51436-51482 FS-3. 3 

## Student Notes

On a Series 700 HFS implementation, a hard disk may be comprised of a LIF directory, a file system area, a swap area, and a boot area. The **root disk** will hold all four of these areas. (The root disk or system disk is the disk that holds the root file system.) At the start of the root disk is a 8KB area reserved for the boot. This 8KB area is a file system in its own right, albeit not a a UNIX file system. Its format is Logical Interchange Format or LIF. LIF is a very simple file system with simple data structures that can easily fit in a small area such as 8KB. Its simplicity makes it a good choice at boot up in order to keep the code size small in the boot ROM. On a Series 700, this LIF area is actually a directory to direct the boot up process to the last 2MB on the root disk that actually holds, the bootstrap utility, the boot up code that will load in the HP-UX kernel, /hp-ux. This 2MB area is actually taken from the last 2MB of swap on the root disk. Swap will occupy the end of this disk save the last 2MB. The file system area is located between the 8KB LIF directory and swap for the root disk.

Non-root disks will typically contain a single swap area, a single file system, or a combination of a both a single swap area and a single file system area. On the Series 700 computer there is only one file system area per physical disk.

## FS-4. SLIDE: Disk Layout (Series 600/800)

### Disk Layout (Series 600/800)

Disks can hold:

- Boot area (2MB)
- File system
- Swap

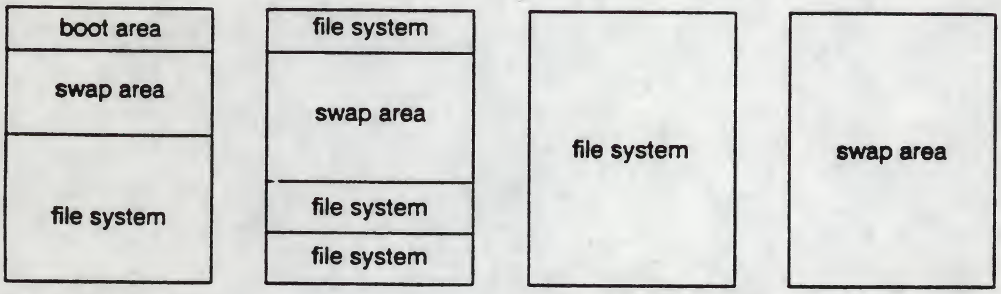


typical root disk

typical non-root disks



## Student Notes

On a Series 600/800 HFS implementation, a hard disk may be comprised of a boot area, multiple file system areas, and a swap area. The **root disk** will hold all three of these types of areas. (The root disk or system disk is the disk that holds the root file system.) Each of these areas, boot, the multiple file systems, and swap, is know as a **disk section** or also known as a **disk partition**.

At the start of the root disk is a 2MB area reserved for the boot. This 2MB area is a file system in its own right, albeit not a a UNIX file system. Its format is Logical Interchange Format or LIF. LIF is a very simple file system with simple data structures that can easily fit in a small area such as 2MB. Its simplicity makes it a good choice at boot up in order to keep the code size small in the boot ROM. This LIF volume actually holds the boot up code that will load in the HP-UX kernel, `/hp-ux`.

On most root disks actually in use, the next section is usually swap although this is not required. This area is usually followed by the file system area, usually occupied by the root file system. The root disk may contain many other file systems or only the root section.



118

Non-root disks can contain a single swap area, a single file system, or a combination of multiple swap areas and multiple file system areas. For performance reasons, multiple swap areas on a single physical disk is not recommended.

## FS-5. SLIDE: File System Layout

78

### File System Layout

A file system is comprised of:

filesystem info

- Primary superblock
- Multiple cylinder groups.

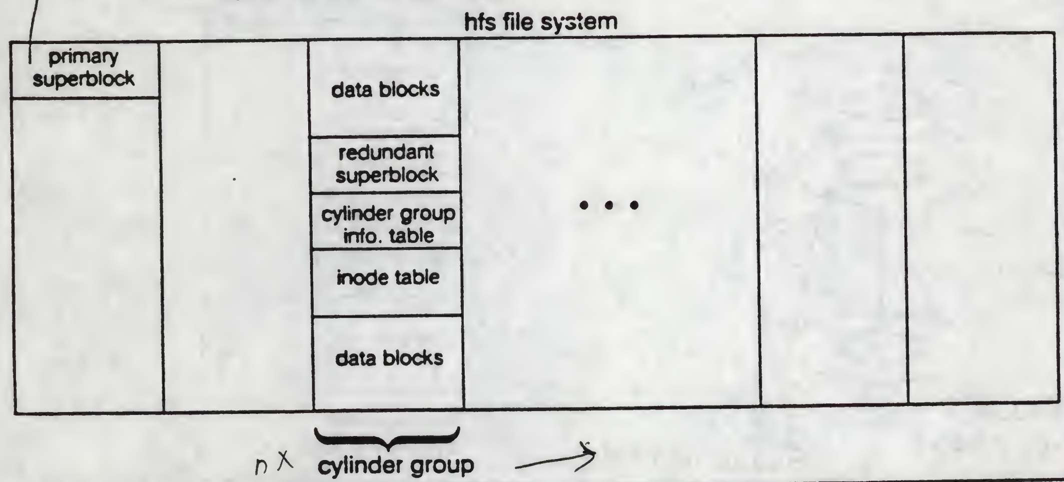


51436-51482 FS-5. 5 

## Student Notes

The layout of the hfs file system is shown on the slide.

The file system is made of:

- Primary superblock
- Multiple cylinder groups.

The **superblock** is a contiguous 8 K of disk space, which HP-UX uses to keep track of the current state of the file system. Each **cylinder group** contains a copy of the primary superblock, a cylinder group information table, an inode table, and data blocks.



119

## FS-6. SLIDE: The Cylinder Group

79

### The Cylinder Group

- Physically contiguous area of disk cylinders
- Contains a copy of the redundant superblock, cylinder group information structure, inode table, and data blocks

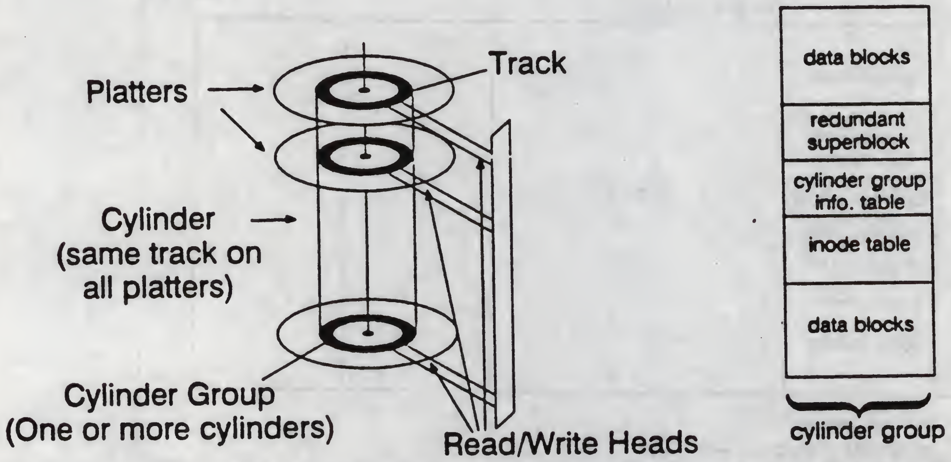


51436-51482 FS-6. 6 

## Student Notes

A **cylinder group** is a group of one or more disk cylinders.

A cylinder is a collection of tracks formed as the head-disk-assembly (HDA) positions all the heads on multiple platters (disk surfaces) at the same distance from the edge of the disk surfaces. The physical picture may help you visualize the concept of a "cylinder" of data being traced out by the array of heads.

Each cylinder group contains a copy of the primary superblock, a cylinder group information structure, an inode table, and data blocks. The cylinder group controls all accesses to a file and its associated data.

We already discussed the superblock. A copy of the superblock is located in each cylinder group so that any single track, cylinder, or platter can be lost without losing all copies of the superblock.

The **cylinder group information** contains the dynamic parameters of the cylinder group:

- Number of inodes and data blocks

- Pointers to the last used block, fragment, and inode
- Number of available fragments
- Used inode map
- Free block map

The **inode table** contains entries for a set of inodes. Inodes contain information about individual files. The number of inodes allocated per cylinder group is determined when the file system is created and cannot be changed once the file system is made. Therefore, the default allocated by HFS is more than will be needed for average usage.

The **data blocks** are the actual data in a file. These are referenced by an addressing scheme in the inode which will be discussed shortly.

Directory information is contained within normal data blocks and does not really occupy a distinct region within a cylinder group. Any unused data block may be allocated for use as a directory. We will talk about directory information a little later.

Cylinder group information is updated when ever the **sync** command is executed.

## FS-7. SLIDE: The Superblock

8 C

### The Superblock

The superblock contains:

- Critical information about the file system
- Static information which includes:
  - File system size
  - Block size
  - Fragment size
  - Disk characteristics
- Dynamic information which includes:
  - Total number of free data blocks
  - Total number of free inodes
  - File system clean flag

51436-51482 FS-7. 7 

## Student Notes

As already mentioned, the superblock is a contiguous, 8KB of disk space which HP-UX uses to keep track of the current state of the file system.

HP-UX uses information in the superblock for various file system maintenance procedures, for example, when you mount a file system or perform a file system check.

Because the superblock is so important, HP-UX always keeps a copy of the superblock in main memory. The superblock on disk is updated whenever the `sync` command is executed. HP-UX also keeps a copy of the primary superblock in each cylinder group. If the primary superblock is lost, damaged, or becomes corrupted in some way, then a copy can be retrieved from one of the cylinder groups. The redundant copies of the superblock contain only the information identified as static on the slide. The dynamic information can be reconstructed with a command called `fsck` that we will examine later.

A list of the locations of the redundant superblocks in each cylinder group can be found in a disk file file created when the filesystem was made. The file name depends upon the type of system you have. On the

Series 300/400/700, this file is `/etc/sbtab`. On the Series 600/800 it is called `/etc/super_blocks`. You should print a copy of this file and save it for future reference.



121

## FS-8. SLIDE: A File's Data Structures

81

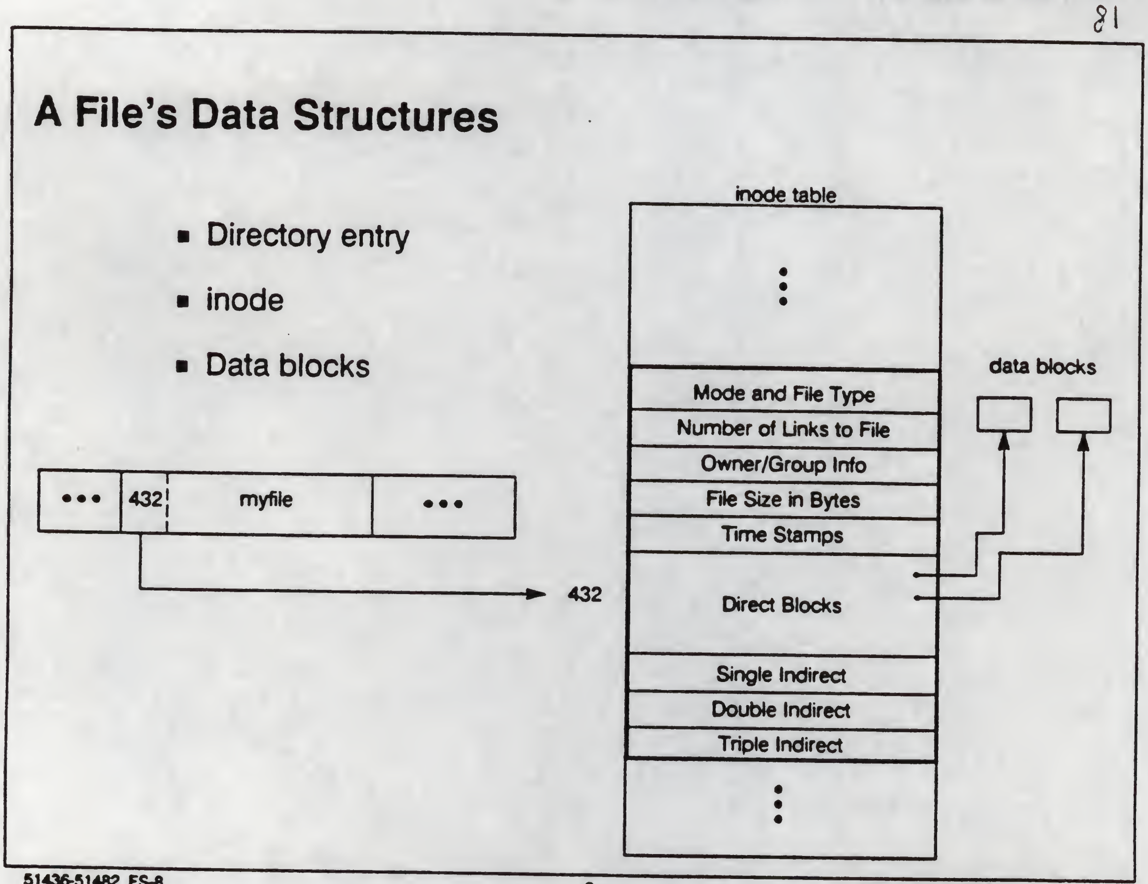



## Student Notes

Files in an HP-UX filesystem are composed of three components:

**Directory entry** The directory entry of a file has two pieces of information, the file's name and its inode number or inum. A directory itself is a file with a particular format, namely multiple directory entries. The directory entry "maps" a particular filename to an inode number. The inode number "maps" to an i-node.

**inode** The inode or information node contains the important information about the file, such as the information output from `ls -l`.

**Data blocks** Another piece of information contained in the inode is the addresses of the data blocks. The data blocks are what actually hold the file's data.

A **directory** behaves like an ordinary file in the file system, except that no user can write into a directory. The function of a directory is to locate all files one level underneath it in the file system tree.

To accomplish this, it contains an entry for each file it must locate. The entry contains the filename and its inode number.

On HP-UX, you can have a file system that supports either short filenames (up to 14 characters) or long filenames (up to 255 characters). The 14-character length is the standard. The following discussion applies to the standard 14 character length.

A directory entry (also known as a directory slot) consists of four fields:

- Binary inode number (4 bytes)
- Length of this directory entry (2 bytes)
- Length of file name (2 bytes)
- File name (up to 14 bytes)

In the example, there is a picture of a directory structure. If we tried to access `myfile`, the corresponding inode number 432 would be retrieved from the directory and used as a pointer to the inode 432 in the inode table. The data block pointer of inode 432 would finally lead to the actual data of `myfile`.

When a directory is first created, it immediately contains two entries—dot (".") and dot dot (".."). These entries are used for relative path addressing. The entry for "." will contain the inode number for this directory itself, while the entry for ".." will contain the inode number of its parent. The only exception occurs when the directory is the top-level directory (root) of the file system. In this case, the inode number will always be 2. Since it has no parent, the number 2 is placed in its "." and ".." entries.

When a file is created, the file name and the inode number are placed in the first slot in the directory that has a 0 in the inode number field. When a file is removed, a 0 is placed in its inode number field. A 0 signifies an empty slot that is available for re-use.

We will come to discover later that there are important advantages to separating the information about a file (its inode) from the contents of the file (its data blocks) and its name (its directory entry). This architecture is the basis for several very nice capabilities for file management. We will cover this later in more detail. This slide should help with the "Big Picture" of files, to see where the "inode" fits into the overall scheme of things in the filesystem.

## FS-9. SLIDE: Inodes

82

### Inodes

- There is one inode for each possible file
- Each inode table entry is 128 bytes long
- A large fixed number of inodes is allocated when the file system is created
- The inodes for a particular file system are distributed among the inode tables of the cylinder groups in that file system
- An inode has a unique inode number, used to locate the cylinder group and offset into the inode table.

51436-51482 FS-9. 9 

## Student Notes

A file is accessed through information stored in its inode. The inode contains every piece of information about an HP-UX file except one. The missing piece is the file name (which is stored separately, in a directory, everything else you would want to know about a file is in the inode. There is one inode for each existing file. Consequently, all files in a file system have an associated inode. To see the inode number for a file, use the `ls -i` *filename* command.

The inodes for each cylinder group are maintained in the inode table (within the cylinder group). Since the inode is the one place in the file system where (almost) all information about a file is stored, the inode table is accessed often.

The contiguous nature of the inode table makes accessing the table simple. Here, contiguous means that all of the information about files in a cylinder group is stored together in one place near the center of the cylinders. The inode table is physically located immediately after the cylinder group information structure on the same cylinder group.

Each inode is referred to by a unique number called the `i number` for the file it represents. The `i number` is used (by the filesystem code) to calculate which cylinder group contains the inode. The calculation also yields its relative position in the inode table. For instance, since an inode is 128 bytes long, the fifth inode in the list will start 512 bytes from the beginning of the inode table Since the cylinder group information structure knows the length of the inode table, it is easy to determine when the end of the list has been reached.

## FS-10. SLIDE: Inode Content and File Mapping Scheme

83

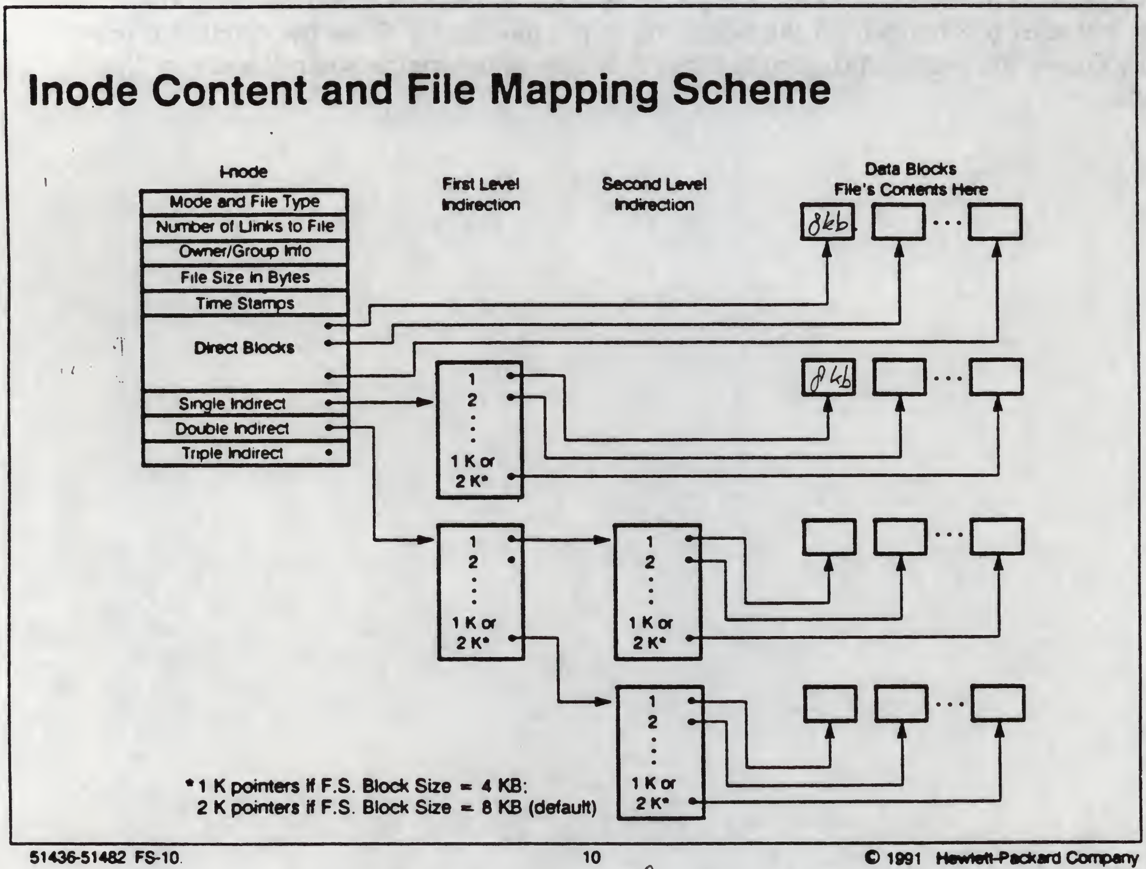


ls -i

## Student Notes

The inode for a file contains almost all the pertinent information "about" a particular file. The primary information contained in an inode is:

- the mode or permissions of the file
- the type of file (that is, regular, directory, special)
- the number of hard links to the file
- the current owner of the file
- the group associated with the file
- the real size of the file (more may be allocated)
- time stamps relating to file activity
  - time/date of last file data change,

- time/date of last file access
- time/date of last inode modification

- disk addresses, or pointers to disk addresses, where the file's data is stored.

Two things that are not kept in the inode are the file's name and its data.

HP-UX implements file-level security through the inode by maintaining file and directory permissions in the inode. Since the inode knows where the file's data is, all accesses to a file take place by first accessing the inode. This puts all the security data directly in-line with any file access. Note that the inode is the only thing that knows where to find a file's data.

The inode contains slots for 15 disk addresses or pointers. The addresses stored in an inode will be a fragment addresses (disk address). (We'll talk about block and fragments on the next slide.) An entry is made in one of these slots when a file is created or when an existing file grows and occupies additional space. If the size of a file is decreased, resulting in the de-allocation of data fragments, the corresponding addresses are removed from the inode as soon as a data block is released.

## FS-11. SLIDE: Block and Fragment Sizes

84

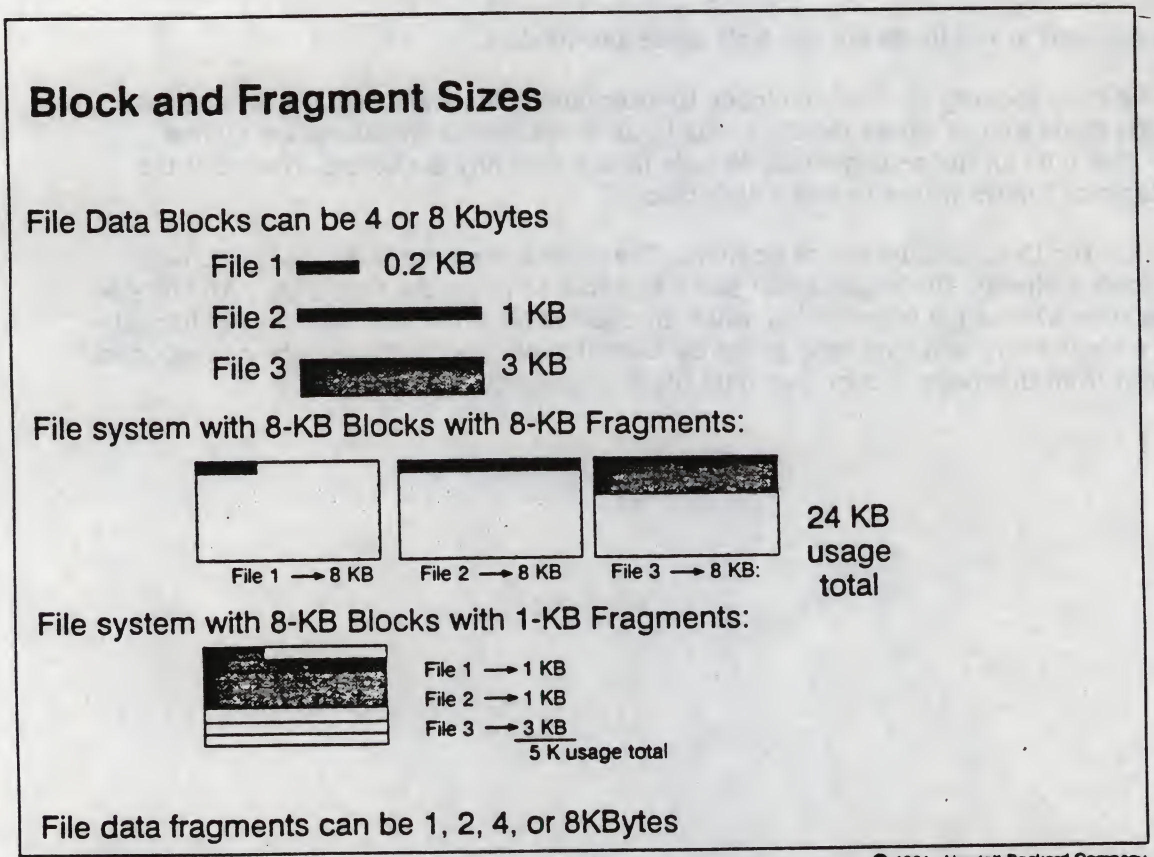


51436-51482 FS-11. 11 

### Student Notes

Each cylinder group contains data blocks in addition to the superblock, the cylinder group information table and the inode table. These areas are used to store the data of regular files, directories and special files. Free space in these areas is allocated in block sizes. **Blocks**, which can be either 4 K or 8 K, form the smallest physical unit that can be accessed by the HP-UX system on disk.

A data block is divided up into 1, 2, 4, or 8 K fragments. The individual fragments within one block form the smallest logical unit that can be accessed by HP-UX. When allocating blocks and fragments for a file:

- A fragment will only be allocated at the end of a file (in other words, full blocks are allocated until less than a full block is left).
- A file will use fragments contained within one block.
- A block may contain fragments from more than one file.
- Multiple fragments may be used by a file, but only at the end of the file, and they must be contiguous within one block.

The block and the fragment sizes are specified at file system creation and cannot be modified without recreating the entire file system. Having large block or fragment sizes has both benefits and costs. If your application has large files, a large block and fragment size could significantly reduce the number of disk accesses, thereby increasing file system throughput. Eight K blocks with 8 K fragments are typically found in data base applications. Large block and fragment sizes are also typical in the file system containing the `/tmp` directory. Here, the files are typically volatile and temporary, so fast access is more important than disk space economy.

If your application uses small files (most HP-UX files are small), then a large block or fragment size wastes space. Accordingly, a combination of a 4 or 8 K block size with a 1 K fragment size is common for file systems containing directories such as `/users`.

### Allocation Rules

1. Each inode will have `n` pointers allocated.
2. Where `n` is calculated using Integer division.

   `n = file-size (bytes) / data-block-size`
3. The first (`n`) pointers point to whole data blocks.
4. If there is a remainder, another pointer is used. It is a fragment pointer that points at the rest of the data.
5. The last pointer (`nth + 1`), if needed, points to the first of enough contiguous fragments to store the rest of the data.
6. There is **at most** one fragment pointer for a file, and it will always be the last pointer for that file.

## FS-12. SLIDE: Review of HP-UX Hard Links

85

### Review of HP-UX Hard Links

- Syntax:

`ln file1 [ file2 ] target`

- Example:

`# ln wing-data test-data`

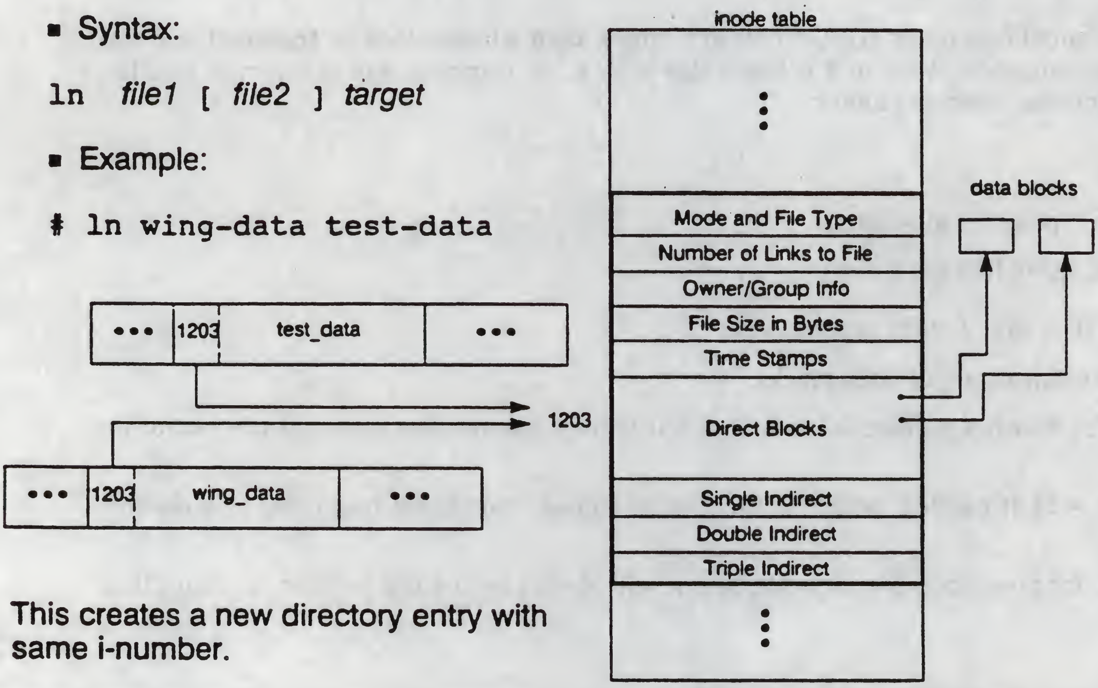


This creates a new directory entry with same i-number.

51436-51482 FS-12. 12 

## Student Notes

- Links are different directory entries indicating the same file
- Links are sometimes called "hard" links
- Links should be used when a file needs two or more different names.
- Example:

  `# ln /users/projectA/src/control/test/data/users/bob/data`

If necessary, hard links can be used to link files (or directories) that do not cross file system boundaries. It is quite useful to be able to link two files together.

Having two different directory entries point to the same file can be a bit confusing, but it is widely used in HP-UX. The "-l" option to the **ls** command is used to see if there are links to a file. If the number in the second field is greater than 1, there are other "names" for that file.

In the first example on the slide, a user application requires that a file be named two different names in the same directory. The file `wing-data` must exist for the command to work. The figure illustrates the state of the file system after the link command finishes.

When the command is executed it creates a directory entry which has the new file name (the target name) associated with the same i-number as the original file. Both directory entries are then said to "point to the same file". A field in the i-node is incremented to indicate the presence of the link. Only after the link count is reduced to zero does the actual data and inode get recycled.



126

## FS-13. SLIDE: Symbolic Links

8 6

### Symbolic Links

- Link directories or files across file systems
- Should be used only when absolutely necessary
- Syntax:

*ln -s file1 [ file2 ] target*

- Example:

```
# ln -s /aux/local_bin /usr/local/bin
```

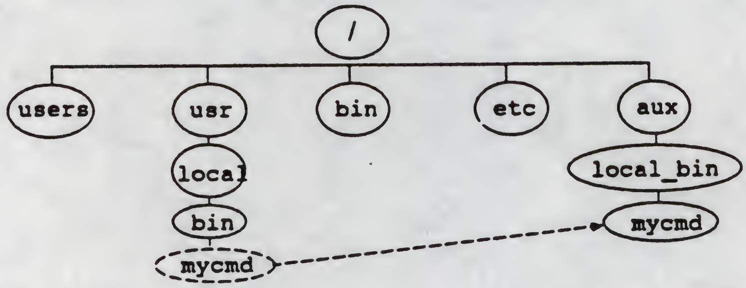


51436-51482 FS-13. 13 

## Student Notes

If necessary, symbolic links can be used to link files (or directories) across file system boundaries. It is quite powerful to be able to link two directories together. This allows directories with a specific path name to be physically put on any filesystem. Linking directories together is helpful on a Series 800 if a particular disk section is not large enough to hold an entire directory structure. Part of the directory structure can then be put on another file system.

Since having two different directories point to the same files can be a bit confusing, it is recommended that symbolic links be used sparingly.

In the example on the slide, a user application requires that several large files be put in the `/usr/local/bin` directory. The system administrator has determined that the filesystem that holds `/usr/local/bin` does not have enough available disk space to hold these files. Therefore, he/she decides to create a symbolic link between the `/usr/local/bin` directory and the `/aux/local_bin` directory, since the available disk space in the file system where `/aux/local_bin` is located is much larger than that of file system where `/usr/local/bin`. The following command is executed:

```
# ln -s /aux/local_bin /usr/local/bin
```

This command creates a directory /aux/local_bin and links it to the directory /usr/local/bin. This means that any files written to the directory /usr/local/bin will physically be put under the directory /aux/local_bin. However, the symbolic link between the two directories will allow logical access from either directory structure. For example:

```
$ cp myfile /usr/local/bin
```

will create a file /aux/local_bin/mycmd However due to the symbolic link, a directory listing of either /users/thewiz/subdir or /tmp/subdir will show the existence of the file.



127

## FS-14. SLIDE: Series 600/800 Disk Sectioning

### Series 600/800 Disk Sectioning

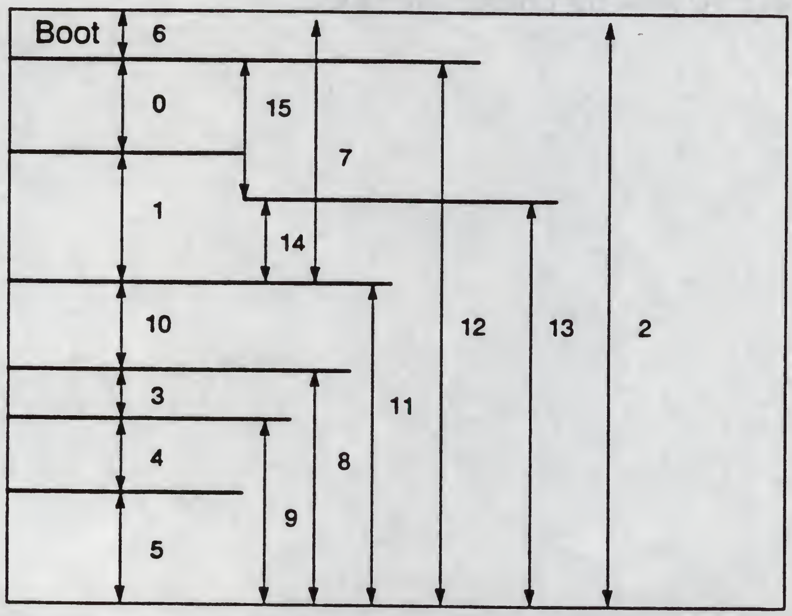


51436-51482 FS-14.    14    

## Student Notes

As mentioned earlier, multiple file systems can exist on a single disk. This is accomplished by sectioning (or partitioning) the disk and creating file systems in the various sections. These are logical sections imposed strictly by the software. Disk sectioning allows the operating system to treat each section as if it were a separate disk.

The advantages of employing a disk sectioning scheme generally outweigh the added complexity of managing such a scheme. Among the chief advantages is the ability to control the amount of disk space that may be used by a certain project or group of users. The administrator can select appropriate sizes based on the application. Since a file system is contained within and limited to the size of the disk section that it is created in, the file system cannot grow without bounds. By the same token, it does not make sense to use two disk sections that overlap. If overlap were to occur, two file systems would be competing for the same disk space which would undoubtedly lead to disaster.

**Warning** **The HP-UX file system will not prevent you from creating a file system in a section that encompasses a section that is already in use!**

The boot area resides on section 6, by default. This implies that the use of section 7 or section 2 on your boot disk is prohibited since section 6 is contained within sections 7 and 2. (Note that section 2 is always the entire device.)

The number of sections that can be defined on a hard disk varies with disk models. The size of each section is also dependent on the disk model. The **/etc/disktab** file defines the section sizes and locations for the various supported disk models. A sample **/etc/disktab** file is shown below.

```
# The recommended layout for file systems is as follows:
# Common for all devices:
#       Section 6       boot section
#       Section 0       root file system
#       Section 1       swap section
#       Section 3       /tmp
# For larger devices (have Section 10):
#       Section 4       /usr
#       Section 5       /extra
#       Section 10      /mnt
hp7937:\
:ty=winchester:ns#30:nt#13:nc#1396:rm#3600:\
:s0#24280:b0#8192:f0#1024:\
:s1#48560:\
:s2#558051:\
:s3#29298:b3#8192:f3#8192:\
:s4#107426:b4#8192:f4#1024:\
:s5#216664:b5#8192:f5#1024:\
:s6#1998:\
:s7#75152:b7#4096:f7#1024:\
:s8#353778:b8#4096:f8#1024:\
:s9#324196:b9#4096:f9#1024:\
:s10#129024:b10#4096:f10#1024:\
:s11#482898:b11#8192:f11#1024:\
:s12#556052:b12#4096:f12#1024:\
:s13#507282:b13#4096:f13#1024:\
:s14#24280:b14#8192:f14#1024:\
:s15#48560:
```

## FS-15. SLIDE: Series 600/800 HP-7937 Disk Sectioning

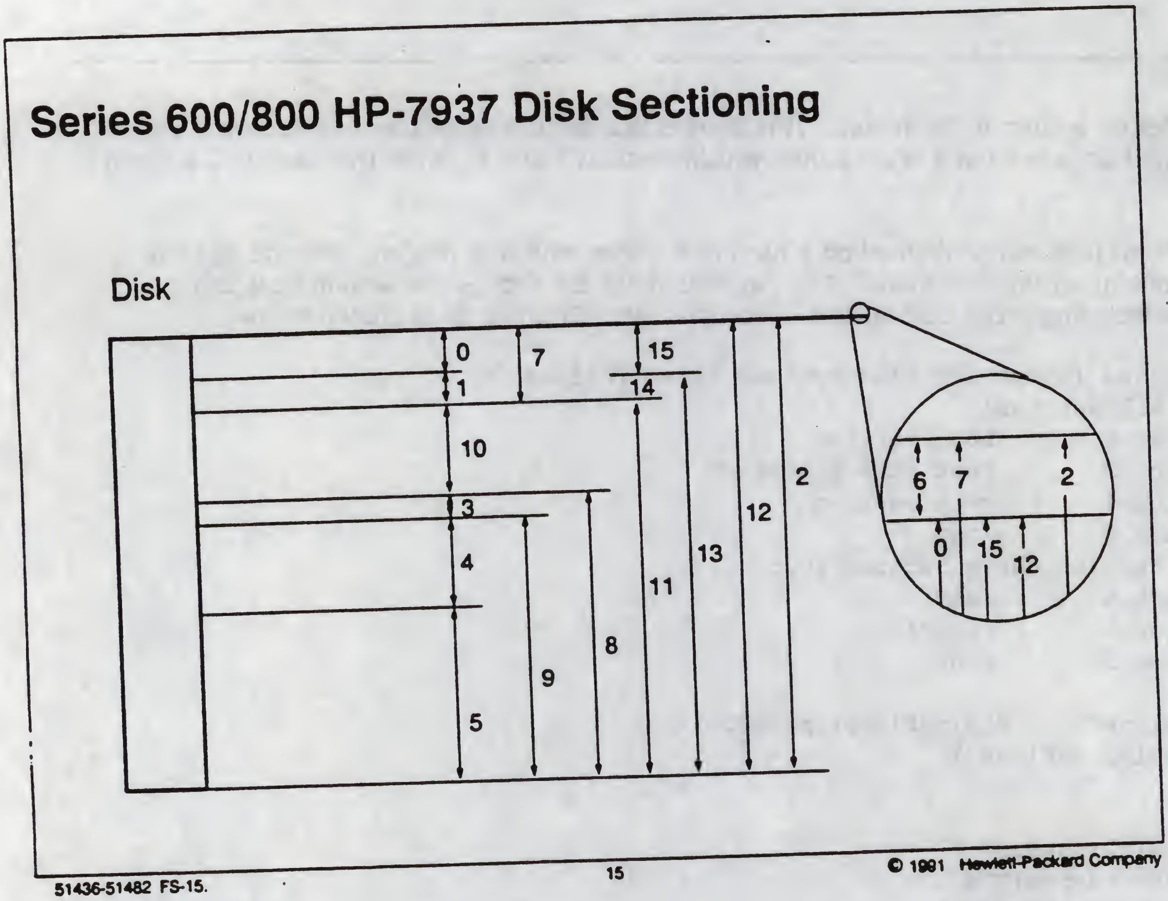


## Student Notes

The slide shows the disk sections for a 7937 disk. Each disk section is drawn in proportion to its size. Because it is drawn in proportion, section 6 (which is only 2 MB) does not show up. Therefore, a small portion of the sections at the very top are blown up so that you can see section 6. (This what you see in on the right hand side of the slide).

The following table shows the section definitions for the HP 7937.

**Table FS-1.**

| Section Number | Capacity in MB | Cylinders |
|---|---|---|
| 0 | 24 | 5 - 65 |
| 1 | 49 | 66 - 187 |
| 2 | 571 | 0 - 1395 |
| 3 | 30 | 511 - 584 |
| 4 | 110 | 585 - 853 |
| 5 | 221 | 854 - 1395 |
| 6 | 2 | 0 - 4 |
| 7 | 76 | 0 - 187 |
| 8 | 362 | 511 - 1395 |
| 9 | 331 | 585 - 1395 |
| 10 | 132 | 188 - 510 |
| 11 | 494 | 188 - 1395 |
| 12 | 569 | 5 - 1395 |
| 13 | 519 | 127 - 1395 |
| 14 | 24 | 127 - 187 |
| 15 | 49 | 5 - 126 |

Please refer to the *Configuration Reference Manual / HP 9000 Series 800 Computers* to determine the disk sectioning scheme of your disk type.

## FS-16. SELF-STUDY QUESTIONS: File System For Series 300/400 & Series 600/800

### Directions

Complete the following questions.

1. List the components of an HFS file system.

2. What is the function of the superblock? What information does the superblock contain? When is this information updated?

3. List the contents of a cylinder group.

4. What is an inode? List at least five pieces of an inode.

5. Where in the file system are the individual file names held?

6. What happens to a directory entry when a file is removed?

7. How is the data for a file stored on disk?

8. In a file system where the block size is 8 K, how large can a file grow before HP- UX must use single indirect data block addressing? How large can the file grow before double indirect data block addressing must be used? (Note that an 8 K data block can hold 2048 data block pointers.)

9. What is the boot area and where does it reside?

10. Series 800 users only: What is disk sectioning?

## FS-17. LAB: Series 300/400 File System

1. Use the `df -t` command to see the number free blocks on the root file system (/). (Note a `df` block is 512 bytes.) Create a file. Use `df -t` again. Did blocks decrease? If so, by how much? Try the `bdf -i` command as well. What are the differences in the output between `df` and `bdf`?

2. Use the `ls -i` command to find the inode numbers associated with the entries in the root directory.

3. Create a file called `test.pattern` with some lines in it. Use the `ls -i` command to find out its inode number. Now, rename the file to your name (use the `mv` command). Do another `ls -i`. Has its inode number changed? Make a copy of your new file into `test.pattern` (use `cp`). Does this file have the same inode number? Now hard link another file (`test2`) to `test.pattern`. What inode number does `test2` have?

## FS-18. LAB: Series 600/800 File System

1. Create a file called `test.pattern` in your home directory. Use the `ls -i` command to find out its inode number. Now, rename the file to your name (use the `mv` command). Do another `ls -i`. Has its inode number changed? Make a copy of your new file into `test.pattern` (use cp). Does this file have the same inode number? Now hard link another file (`test2`) to `test.pattern`. What inode number does `test2` have?

2. What is the inode number associated with /hp-ux?

3. Often the `/tmp` directory resides on a file system other than the `/users` directory. Try to link a file in your home directory to a file name in `/tmp`. What happens? Why?



130B

# Module CF — File System Creation (Series 300/400/700)

## Objectives

Upon completion of this module, the student will be able to:

- Add or remove a disk drive.
- Initialize media in HP-UX format.
- Create a file system.
- Mount or unmount a file system.
- Automatically mount a file system.
- Determine free disk space.
- Add or remove a CD-ROM file system.
- Use SAM to help manage a file system.
- Manage disk quotas in the file system



131

## CF-1. SLIDE: Creating a New File System

?7

### Creating a New File System

1. Install the hardware
2. Create device files for the new device (`mknod`)
3. If necessary, initialize the media (`mediainit`)
4. Make the file system (`newfs`)
5. Make the file system available for use (`mount`)
6. Add the new file system to `/etc/checklist` for automatic mounting

51436 CF-1. 1 

## Student Notes

The process for creating a file system is through the use of a series of commands that you may enter interactively, or by invoking the SAM utility. A Series 300/400/700 system only allows one section per hard disk, so, one hard disk drive can contain only one file system.

There are many reasons why you may want to add a new file system. Some reasons include:

- You anticipate that your current file system will soon reach maximum capacity.
- Your current file system has already reached maximum capacity.
- You wish to physically separate portions of a file system.

To create a new file system, just follow the steps on the slide. We will talk about each step in detail on the following pages.

If you are creating a new file system on a new disk drive, you must connect the device to the system first. For help in this area, read the Installation Manual for the device that you are connecting.

# Module CF — File System Creation (Series 300/400/700)

Theoretically, many devices can hold file system; however, we recommend that you use a hard disk to hold an HP-UX file system. Flexible disks can be used to hold a file system, but they have limited capacity and are slow. Cartridge and mag tape drives should not be used to hold a file system. They have very slow access times and the tape has a short life span due to wear and tear. Rewritable magneto-optical (MO) disks are typically used to backup a file system, but can be used as an auxiliary file system. They can be mounted as a filesystem to facilitate transparent archival and retrieval of files. The MO disks are slower than hard disks, but substantially faster than flexible disk or tape.



133

## CF-2. SLIDE: Review: Creating Device Files For the New Hard Disk

8cP

**Review:**

**Creating Device Files For the New Hard Disk**

Series 700 examples:

```
mknod /dev/[r]dsk/BsO c|b major 0xSSFBUV
mknod /dev/rdsk/4s0 c 47 0x201400
mknod /dev/dsk/4s0 b 7 0x201400
```

Series 300/400 examples:

```
mknod /dev/[r]dsk/BsO c|b major 0xScBaUV
mknod /dev/rdsk/4s0 c 4 0x0e0400
mknod /dev/dsk/4s0 b 0 0x0e0400
```

| | HP-IB | SCSI |
|---|---|---|
| character | 4 | 47 |
| block | 0 | 7 |

51436 CF-2. 2 

### Student Notes

The Series 700 minor number format is 0x*SSFBUV* where:

***S*** System bus module number
SCSI SBM numbers can be:

- 2 Core I/O board
- 4 EISA Bus Adapter

***S*** EISA slot number
EISA slot number is only used when for cards located off of the EISA bus. The SBM is 4 for the EISA bus adapter.

***F*** Function number
For the core input/output board the SBM is 2 and the SCSI function number is 1.

***B*** SCSI bus address number.

*U* Unit number
The unit number is typically 0. A dual-drive floppy disk unit is included in the exceptions to this generalization. The left drive is unit 0 and the right is unit 1.

*V* The volume number is always 0.

The Series 300/400 minor number format is 0x*ScBaUV* where:

*Sc* Select code
For SCSI and high-speed HP-IB the select code is 14 or 0x0e. For standard-speed HP-IB the select code is 7 or 0x07.

*Ba* SCSI or HP-IB bus address number on a Series 300/400

*U* Unit number
The unit number is typically 0. A dual-drive floppy disk unit is included in the exceptions to this generalization. The left drive is unit 0 and the right is unit 1.

*V* The volume number is always 0.

See the module on creating Device Files for a more detailed review.

You must create both character and block device files for your disk drive. Character device files are used to create the file system with the `newfs` command, while block device files are used with the `mount` and `unmount` commands and also used to access the file system.

## CF-3. SLIDE: Initializing the Device

89

### Initializing the Device

- Use the `mediainit` command to initialize a device.
- `mediainit:`
  - Formats the media
  - Performs read/write tests to verify integrity
  - Marks as unusable blocks that fail the tests
  - Assigns an interleave factor (if necessary)
- Syntax:

`mediainit [-vr] [-f` *format*`] [-i` *interleave*`]` *char_dev_file*

```
# mediainit /dev/rdsk/2s0
```

```
# mediainit -v -f 3 -i 2 /dev/rdsk/3s0
```

51436 CF-3. 3 

### Student Notes

The `mediainit` command is a command to re-format the disk or tape, rather than re-initialize it as its name might imply.

The `mediainit` command "initializes" mass storage media by formatting the media, writing and reading test patterns to verify media integrity, then sparing any defective blocks found. This process prepares the disk or tape for error-free operation. After initialization, the media can contain data, files, or a file system.

All Hewlett Packard hard disks are initialized at the factory before shipment. You must initialize you hard disk if:

- your disk is known to be damaged, or
- you want a different interleave than what the disk currently has

Many of the SCSI disks require a very specific interleave factor for proper media initialization. If you are not certain that the interleave factor is correct for you destination disk, then you should not initialize your disk. It is recommended, though, that you initialize a optical disk or flexible disk before using it the first time.

The amount of time to initialize the disk depends upon the type and size of your disk (intervals of time run from 10 or 20 minutes for high performance disks to 45 minutes for much slower disks). The initialization destroys all data previously existing on that disk, verifies the storage integrity, and marks any defective blocks so they will not be used.

The steps to initialize your media follow:

1. Log in as root.
2. If necessary, turn on the drive. Wait until it completes self-tests and is ready for use.
3. If necessary, insert the media and wait until it is ready for use.
4. Create the character device file for the media you are going to initialize (if it doesn't already exist).
5. Initialize the media with `mediainit` using the character device file.

If this is not a new device, and you are re-initializing it, make sure the device is not mounted. Use `bdf` to make sure that it is not mounted. (We will look at the output of `bdf` on a later slide.)

Some of the options to `mediainit` include:

**-v** Verbose. Normally, `mediainit` provides only fatal error messages directed to diagnostic output (stderr). The `-v` option sends device-specific information related to low-level operation of `mediainit` to standard output (stdout).

**-r** Re-certify. Forces a complete tape certification whether or not the tape has been certified previously. This can take an extra-ordinary amount of time, several hours per tape.

**-i** *interleave* Interleave factor. Refers to the relationship between sequential logical records and sequential physical records. The choice of interleave factor can have a substantial impact on disk performance. If a disk being initialized requires an interleave factor but none is specified, `mediainit` provides an appropriate, though not necessarily optimum default. For CS/80 and SS/80 disks, `mediainit` uses whatever the device reports as its current interleave factor. This option is best defaulted, except for flexible disks. See `/etc/disktab` for recommendations regarding flexible disks.

**-f** *format* Format option. A device-specific number in the range 0-239. It is intended solely for use with certain SS/80 devices that support multiple media formats (independent from interleave factor). For example, certain microfloppy drives support 256, 512, and 1024-byte sectors. This option is best defaulted, except for flexible disks. See `/etc/disktab` for recommendations regarding flexible disks.

*char_dev_file* Path name to the character device file associated with the device to be initialized. You must supply this option. `mediainit` aborts if you lack either read or write permission to the device file, or if the device is currently open for any other processes. It will also abort if the driver for that type of device is not configured into the kernel. This prevents accidental initialization of the root device or any mounted volume.

**mediainit destroys all existing data in the area being initialized. Perform a back up on any existing data before using mediainit.**

## CF-4. SLIDE: Creating a File System

90

### Creating a File System

- Use the newfs command to create a file system
  - Make sure the media is not mounted
  - Make sure an entry for the disk exists in /etc/disktab
- Syntax:

```
/etc/newfs [-L|-S] [-n] [-v] [ -mkfs-options ] char_dev_file disk_type

# newfs -b 8192 -f 2048 -m 20 -i 4096 /dev/rdsk/5s0 hp2213A
```

51436 CF-4. 4 

## Student Notes

**newfs** is the command for creating a "new" file system. The **newfs** command is actually creating the necessary infrastructure on the media to support a file system's data structures. Those data structures familiar to the user include files and directories. You can think of the **newfs** command as creating an empty file system to be filled with files and directories later.

**newfs** will use the file /etc/disktab to extract the values needed to create the file system. For example, if you are creating your file system on an HP 7937 disk drive, you will supply the label for the HP 7937 disk from the /etc/disktab file as the *disk_type* option to **newfs**. The parameters in /etc/disktab will be used to create your file system unless you override them with command line options and parameters to the **newfs** command..

The syntax of the command is described in more detail below:

**/etc/newfs** [-L|-S] [-n] [-v] [ *-mkfs_options* ] *char_dev_file disk_type*



136

[-L|-S] HP-UX file systems support long file names, 255 characters max, or short file names, 14 characters max.

[-n] Prevent the bootstrap programs from being installed. On a Series 700, **newfs** does not install the bootstrap programs in order to save space.

[-v] Verbose

*mkfs_options* include:

-s *size* Specifies the total size of the file system in blocks. If this option is not specified, **newfs** will calculate the maximum size file system possible on the disk being used.

-b *block_size* Specifies the primary block size in bytes of a file in the file system . Valid file system block sizes are 4096, 8192, 16384, 32768, and 65536. If defaulted, this is 8 Kbytes.

-f *frag_size* Specifies the fragment size of the file system in bytes. The fragment size represents the smallest amount of disk space to be allocated to a file. It must be a power of two not smaller than 1024 (DEV_BSIZE) and no smaller than one-eighth of the file system block size. Default is 1024 bytes.

-m % *minfree* This indicates the percentage of free space reserved. If the amount of free space in the file system falls below this percentage, the superuser is the only person who can write to the file system. This free space is actually reserved to help ensure good performance of the data block allocation policies. Making this parameter less than 10% will impact file system performance negatively. Default is 10 percent.

-i *bytes/inode* This specifies the number of inodes that will be created. In other words for every $x$ number of bytes, one inode will be created. If defaulted, this is 1 inode per 2048 bytes of file space. This is usually more than what is actually needed since 1 inode per 2048 bytes implies an *average* file size of 2048 bytes. If fewer inodes are desired, a larger number should be used.

*char_dev_file* This is the character device file for the device on which you are creating the file system. It must be supplied.

*disk_type* This is the label into the `/etc/disktab` file. It must be supplied.

Typically, the default values for the disk being used are adequate for file system purposes. There may be an occasion, however, when it's necessary to specify a different value. In such a case, the options to **newfs** are very useful. If an option is not specified, **newfs** will either calculate an appropriate value or extract a default value from `/etc/disktab`.

Before using **newfs**:

- Check to see that the media is not mounted. Use the `bdf` command.
- Find an entry to `/etc/disktab` for the drive that will contain the file system. If an entry does not already exist, you must create one. The best way to do this is to copy an existing entry and edit it.

The example on the slide is described below:

```
# newfs -b 8192 -f 2048 -m 20 -i 4096 /dev/rdsk/5s0 hp2213A
```

Creates a new file system on a HP C2213A SCSI disk. The device file is `/dev/rdsk/5s0`, the label into `/etc/disktab` is `hp2213A`. The remaining options provide overrides to the corresponding values in

the disktab entry for this type of disk. The block size is 8192, the fragment size is 2048, there is 20% reserved space, and 4096 bytes of file system space per inode.

For a Series 300/400/700 file system, **newfs** appends the superblock locations to the file **/etc/sbtab**. Print a copy of the appropriate file after you create a new file system.

## CF-5. SLIDE: Fields in the /etc/disktab File

91

### Fields in the /etc/disktab File

| | |
|---|---|
| ty | Information about the disk (informational only) |
| ns | Number of 1k sectors per track |
| nt | Number of tracks per cylinder |
| nc | Total number of cylinders in the file system |
| s0 | Size of file system in 1k blocks |
| b0 | Block size in bytes for file system |
| f0 | Fragment size in bytes for file system |
| se | Number of bytes per physical sector (informational only) |
| rm | Revolution per minute |

51436 CF-5. 5 

### Student Notes

We have taken several typical disk tab entries for a variety of disks, to illustrate how entries in the /etc/disktab file may appear.

```
##################################################
# The HP2213A is a SCSI Coyote II
# Total formatted capacity:  663 Mbytes
#   512 Bytes/sector
#   56 sectors/track;  16 heads; 1447 cylinders;
# Total: 648256 1k sectors
hp2213A|HP_2213A|hp660S|hp97548S|HP_97548S:\
 :64 MBytes reserved for swap & boot:ns#28:nt#16:nc#1302:\
 :s0#583296:b0#8192:f0#1024:\
 :se#512:rm#4002:
hp2213A_300MB|hp660S_300MB:\
        :300 MBbytes reserved for swap & boot:ns#28:nt#16:nc#763:\
```

```
        :s0#341824:b0#8192:f0#1024:\
        :se#512:rm#4002:
hp2213A_96MB|hp660S_96MB:\
 :96 MBytes reserved for swap & boot:ns#28:nt#16:nc#1229:\
 :s0#550592:b0#8192:f0#1024:\
 :se#512:rm#4002:
hp2213A_42MB|hp660S_42MB:\
 :42 MBytes reserved for swap & boot:ns#28:nt#16:nc#1351:\
 :s0#605248:b0#8192:f0#1024:\
 :se#512:rm#4002:
HP_2213A_noswap|hp2213A_noreserve|hp2213A_noswap|hp660S_noreserve|hp660S_noswap:\
 :no swap or boot:ns#28:nt#16:nc#1447:\
 :s0#648256:b0#8192:f0#1024:\
 :se#512:rm#4002:
```

Note that on the Series 700 the boot area is taken from swap space. If you are making the device bootable, you should be aware that there will actually be 2MB less swap available on the device than expected.

Flexible disk examples:

Notice that flexible disks never have swap space, for obvious reasons. They should also not need a boot area.

```
hp9122|hp9122_noswap|hp9122_F3_noswap:\
        :format option 3:ns#5:nt#2:nc#77:\
        :s0#770:b0#8192:f0#1024:\
        :se#1024:rm#600:
```



138

## CF-6. SLIDE: Mounting the New File System

92

### Mounting the New File System

- Create a sub-directory with the **mkdir** command
- Mount the newly created file system with the **mount** command
- Syntax:

`mount` *block_device_file directory*

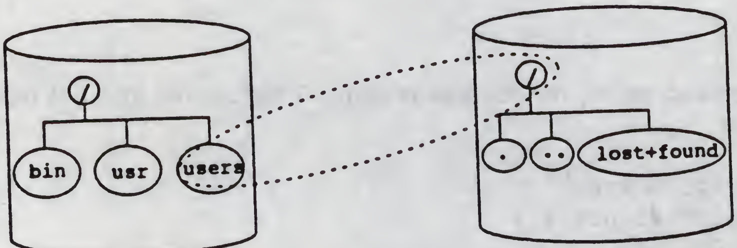


Root File System

New File System
/dev/dsk/2s0

51436 CF-6. 6 

## Student Notes

HP-UX cannot use a file system unless it is mounted. After the file system has been created on the device, it must be incorporated into the root file system. This is done by using the **mount** command to logically associate the root (/) directory on the new file system with a directory on the root file system, the **mount point**.

In the example:

```
# mount /dev/dsk/2s0 /users
```

The new file system whose device file is `/dev/dsk/2s0` is mounted onto `/users` in the root file system.

After the file system is mounted, any reference to `/users` is the same as `/` on the new file system. For example, the following command puts a copy of `/etc/profile` into the root directory of the new file system.

```
# cp /etc/profile /users
```

The mount command with no arguments displays a list of the currently mounted file systems and their respective mount points.

Example:

```
# mount|
/ on /dev/dsk/0s0 read/write on Thu Sep 28 18:16:01 1989
/users on /dev/dsk/2s0 read/write on Fri Sep 29 14:18:00 1989
```

Some options to mount include:

| | |
|---|---|
| -r | mount "read only" |
| -f | "forceable" mount |
| -a | mount "all" file systems in /etc/checklist |

a mount point is busy, you will not be able to mount your new file system on that mount point. A busy mount point is one that is open. For example, if a user performed a cd to the directory that is your mount point, then the mount point is busy. In addition, any files that exist in the directory that is your mount point will be hidden after you mount a file system onto that directory.

You can chain together mounted file systems with certain restrictions when other types of file systems are used. We have only talked here about HFS file systems; there are others. HFS file systems are local file systems on read/write media. CDFS file systems are local file systems on CD-ROM (read only) media. NFS file systems are locally to a remote system and exported and mounted to your local system. The CDFS file system will be covered later, and the NFS file system is covered in the Network Administration course.

You should make certain that the /etc, /dev, and /bin directories never reside anywhere but on the root file system. We already know that the root filesystem is automatically "pseudo mounted" at boot time, so it is assured that these directories will be available when root is mounted. Do not place these on an explicitly mounted file system. File systems other that root are considered explicitly mounted file systems, because they can be "un-mounted" from the system.

**Note** The mount command uses the block device file, whereas mediainit and newfs use the character device file.



139

## CF-7. SLIDE: The umount Command

93

### The umount Command

Syntax: `umount` *block-dev-file*

`umount` *mount-point*

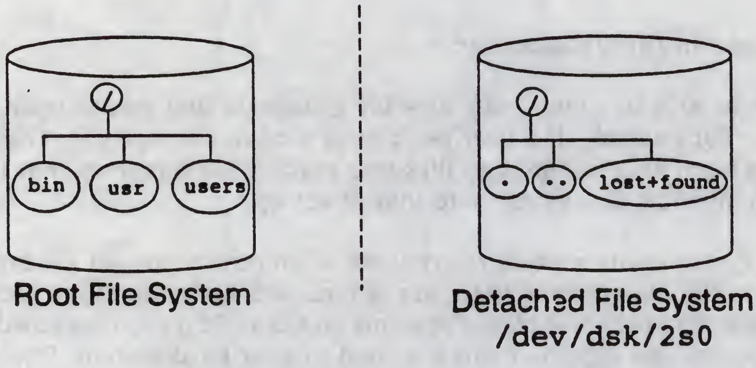


51436 CF-7. 7 

## Student Notes

Now that you know how to mount a new file system, you should also be aware of how to logically disassociate, or unmount, the new file system from the root file system. The command used to unmount the file system is `umount`.

A file system cannot be unmounted if the files system is busy meaning that files are open or if a user's current working directory is a directory in that file system. You can use the `fuser` command to identify which processes are using a file or file structure. Here's some common applications of the `fuser` command.

- `fuser -u /dev/dsk/2s0`
  Lists process IDs and login names of processes using `/dev/dsk/2s0`.
- `fuser -u /etc/passwd`
  Lists process IDs and login names of processes that have the `passwd` file open.

- **fuser -ku /dev/dsk/2s0**
  Terminates all processes that are preventing disk drive 2s0 from being unmounted, listing the process ID and login name of each as it is killed.

Always `umount` all mounted file systems *before* bringing the system down or you may cause corruption to the file systems. The `umount -a` command will unmount all the filesystems in `/etc/checklist`. The `shutdown` script unmounts all file systems before bringing the system down.

You cannot `umount` the root file system.

You cannot `umount` a file system that has dynamic swap enabled on that disk.



140

## CF-8. SLIDE: Example—Moving /users to a New Disk

94

### Example – Moving /users to a New Disk

```
# newfs -n /dev/rdsk/2s0 hp7959S
# mkdir /users.new
# mount /dev/dsk/2s0 /users.new
# cd /users
# find . -print | cpio -pdumv /users.new
# cd /
# rm -rf /users/*
# umount /users.new    (or umount /dev/dsk/2s0)
# mount /dev/dsk/2s0 /users
# rmdir /users.new
```

51436 CF-8. 8 

## Student Notes

Many times when a file system becomes full, the system administrator will want to move the **/users** part of the file system to a new disk. The steps to perform this task are shown on the slide.

You may not be familiar with the `cpio -p` command. In this example, the `cpio -p` command is used to copy a sub-tree of the file system to the new disk. The commands here allow you to have a cookbook approach to solve this problem. You should simply concentrate on the usage of the **newfs**, **mount**, and **umount** commands as a part of the process.

## CF-9. SLIDE: Automatically Mounting File Systems

95

### Automatically Mounting File Systems

- Place an entry in the `/etc/checklist` file
- File systems will be mounted at boot up or you can use the `mount -a` command
- The fields of `/etc/checklist` are shown below:

| block | directory | type | options | backup-frequency | pass-number | #comment |
|---|---|---|---|---|---|---|
| file system (device file) | mount point or attach point | cdfs, hfs, nfs, swap, swapfs, ignore | rw, ro, suid, quotas nosuid, defaults (rw and suid) | 0 | 1,2, . . . etc | |

51436 CF-9. 9 

## Student Notes

Placing an entry in `/etc/checklist` causes your file system to be mounted every time the system boots. The `/etc/rc` script contains the command `mount -a`. `mount -a` mounts those file systems located in the `/etc/checklist` file. The `/etc/checklist` file is *not* maintained by the system and should be edited periodically by the system administrator.

An `/etc/checklist` example:

```
/dev/dsk/0s0    /      hfs  defaults    0 1 # root volume
/dev/dsk/2s0    /users hfs  defaults    0 2 # users
```

The fields of `/etc/checklist` are described in more detail below:

- *block* is the block device file that corresponds to the mounted file system.
- *directory* is the directory to which `mount` mounts the device



141

- *type* is the file system type. Types include:

| | |
|---|---|
| `cdfs` | - local CD-ROM file system. |
| `hfs` | - high-performance (McKusick) file system. |
| `nfs` | - network or remote file system. |
| `swap` | - the device file name is made available as a piece of swap space by the `swapon` command. |
| `swapfs` | - the file system which *directory* resides in is made available as swap space by the `swapon` command. |
| `ignore` | - ignores the entry. Used for comments. |

- *options* is a comma-delimited list of options used by `mount`. Examples are:

| | |
|---|---|
| `rw` | - read/write (default) |
| `ro` | - read only |
| `suid` | - set user-id programs allowed (default) |
| `nosuid` | - no set user-id programs allowed |
| `defaults` | - rw and suid |

- *backup-frequency* is reserved for possible use by future backup utilities.
- *pass-number* is used by the `fsck` command to determine the order in which file system checks are done.
- *comment* is a comment field (must be preceded by a #)

Notice that you shouldn't put primary swap in your `/etc/checklist` file. It causes boot time errors. Only list secondary swap partitions. Details on swap are covered in the swap module of this course.

See the `checklist(4)` in the *HP-UX Reference* manual for more information.

## CF-10. SLIDE: Checking Free Disk Space

96

### Checking Free Disk Space

Syntax:

bdf [ -i ] [ *filesystem* | *file* ]

Examples:

```
# bdf
Filesystem     kbytes     used    avail  capacity    Mounted on
/dev/dsk/0s0   100047    84736    5306      94%      /
/dev/dsk/40s   483392   483392    0        100%      /cdrom


# bdf -i
Filesystem    kbytes    used    avail  %cap iused   ifree  iused Mounted on
/dev/dsk/0s0 4100047   84736    5306   94%  4215    43273  9%    /
```

51436 CF-10. 10 

## Student Notes

The system administrator is responsible for monitoring the amount of free disk space on the system. The easiest way to do this is with the bdf command. The fields have the following meaning:

| | |
|---|---|
| filesystem | Block device file of the file system |
| kbytes | The number of kilobtyes of total disk space on the filesystem |
| used | The number of kilobytes of disk space used by existing files |
| avail | The number of kilobytes of available disk space on the filesystem |
| capacity | The percentage of disk space used by files and the minfree parameter (10% by default for minfree) |
| Mounted on | Directory to which the indicated filesystem is mounted |
| iused | Number of inodes currently in use on the filesystem |
| ifree | Number of free inodes on the filesystem |



142

iused Percentage of inodes used on the filesystem

The -i option adds three columns to the output which give information relating to the availability of inodes in the file system.

## CF-11. SLIDE: CD-ROM File Systems: `cdfs`

### `cdfs`: CD-ROM File Systems

- Allows mass distribution and easy retrieval of large amounts of information
- Read only media
- To use CDFS volumes:

1. Configure the appropriate driver into the kernel
2. Create the necessary device files
3. Mount the CDFS volume with the `mount` command

51436 CF-11. 11 

## Student Notes

CD-ROM is an acronym for Compact Disk-Read Only Memory. Compact Disks (or CDs) contain approximately 600 Mb of data per disk. The information on the CD is virtually permanent; you can read data from a CD, but you cannot write to it. Data on a CD is prepared and mastered using a specialized publishing process. The steps to use a CDFS volume are shown on the slide.

For a Series 300/400 system you can have either an HP-IB or SCSI CD-ROM drive. For a Series 700, only SCSI support is provided. To configure the kernel, add the `cdfs` keyword into the appropriate `dfile`. If you have a SCSI CD-ROM drive, the SCSI driver is needed for the CD-ROM drive, add the `scsi` keyword as well. If you have an HP-IB CD-ROM drive, the HP-IB driver is needed for the CD-ROM drive, add the `cs80` keyword.

Reconfigure the kernel using `config` and `make`, or `sam`.

Once the kernel has been configured, create the necessary device files as you would normally create device files. The device should be created following the conventions described earlier. Then, mount the CDFS volume with the `mount` command.

```
# mount /dev/dsk/5s0 /cdrom-fs
```

Once mounted, HP-UX commands may be used to list, edit (without modifying, as in `view`), print, copy, etc. files on the CDFS volume. In short, anything that does not require modifying the data on the CDFS volumes may be done.

A sample entry in the `/etc/checklist` file for automatically mounting a cdfs volume:

```
/dev/dsk/c4 /cdrom cdfs ro 0 0 # Local CD-ROM drive
```

There are a few commands that are not applicable to CDFS volumes, either because of the read-only nature of the medium, or because they are inherently HFS-oriented. These commands are unsupported under CDFS, and include the following (not a complete list):

- `fsck`
- `fsclean`
- `mediainit`
- `mkfs`
- `ncheck`
- `newfs`

For more information, see the manual: *How HP-UX Works, Concepts for the System Administrator.*

## CF-12. SLIDE: SAM Versus Manual Commands

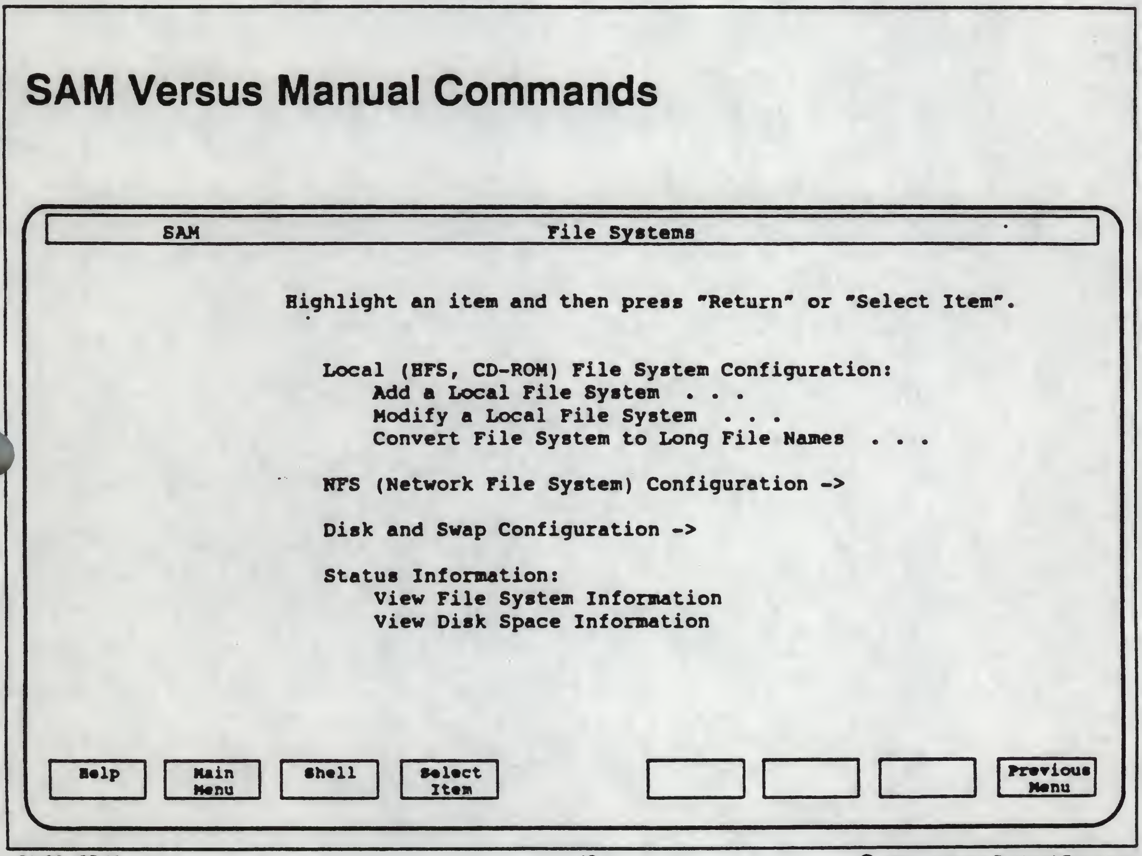


51436 CF-12. 12 

## Student Notes

SAM can be used as well as manual commands to perform file system operations. The tasks that SAM can perform are shown on the slide. This menu will come up after choosing:

```
SAM Main Menu
     |
     V
File Systems ->
```

Before the screen shown in the slide is enabled for interaction, a flashing message will appear at the bottom of the screen:

```
-- working --
```

when SAM is ready for interaction, SAM will highlight the first choice as shown on the slide, and remove the "working" message..



144

# Module CF — File System Creation (Series 300/400/700)

The "Add a Local File System ... " option allows you to add a new file system to your configuration. This may be a hard disk file system, Optical disk (MO) file system, or a CD-ROM file system. It assumes you already have a file system on the disk of the type you are adding. We will visit adding a new disk next.

## CF-13. SLIDE: SAM:Adding a New Hard Disk Drive

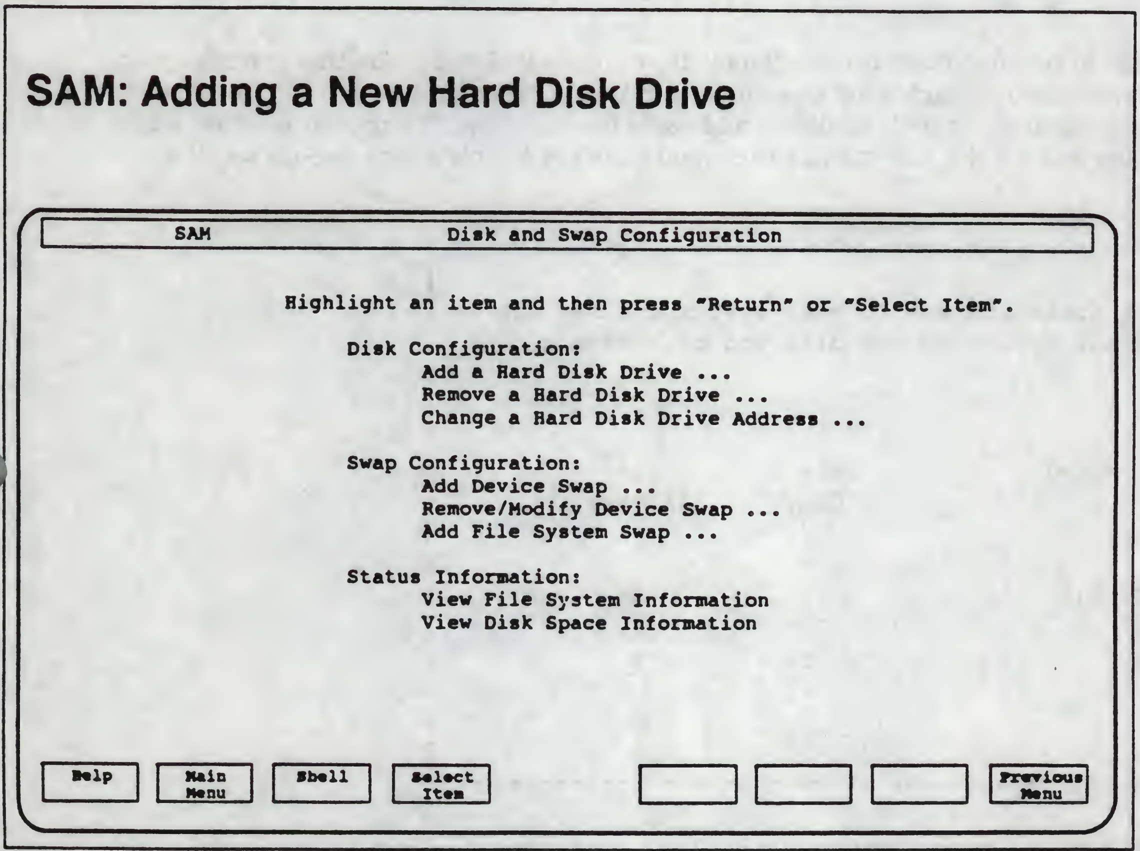


51436 CF-13. 13 

## Student Notes

You will see this screen if you selected "Disk and Swap Configuration" option in the "File Systems" menu. You should take note of the path to follow to reach this SAM menu. We will repeat the path here:

```
SAM Main Menu
     |
     V
"File Systems ->"
     |
     V
"Disk and Swap Configuration ->"
     |
     V
"Add a Hard Disk Drive ..."
```



145

# Module CF — File System Creation (Series 300/400/700)

"Add a Hard Disk Drive" adds a new disk to your configuration and optionally creates a new file system on that disk. Before you see the next menu, there will be a "working" indicator at the bottom of the screen, followed by a pop-up window listing all the candidate disks that SAM could locate on your system.

SAM will display the list of possible disks using a "label" from the "/etc/disktab" file that matches your disk type, which may not exactly match what you think SAM should call your disk drive. The table of candidate disks is displayed along with hardware addresses for each one. The pop-up window will overlay a new menu (we will see the new menu on the next slide.) A sample of the pop-up window might look like:

```
+----------------------------------------------------------+
|                                                          |
|  The following disks are new to your system and may now be |
|  added.  Move the cursor to the disk you wish to add, and |
|  press "Done".                                            |
|                                                          |
|                                                          |
|             Model              Select        Bus         |
|                                 Code       Address       |
|                                                          |
|                                                          |
|        CDROM-HPIB                7            0          |
|        hp7957                   14            2          |
|        hp7959S                  15            1          |
|                                                          |
|                                                          |
|                                                          |
+----------------------------------------------------------+
```

You should choose the disk you want to add with the arrow keys and then press the `Done` softkey. This will allow SAM to partially fill in the menu displayed on the next slide. You can use the (Return) key as a substitute for "Done" if you want.

If your disk is not shown, you can select `Exit Window`, get back to the menu in the slide, where you can exit to a shell. If your disk is not displayed in the list, you should follow precautions described in the installation module to determine that your disk is correctly connected and configured.

## CF-14. SLIDE: SAM: Add a Hard Disk Drive

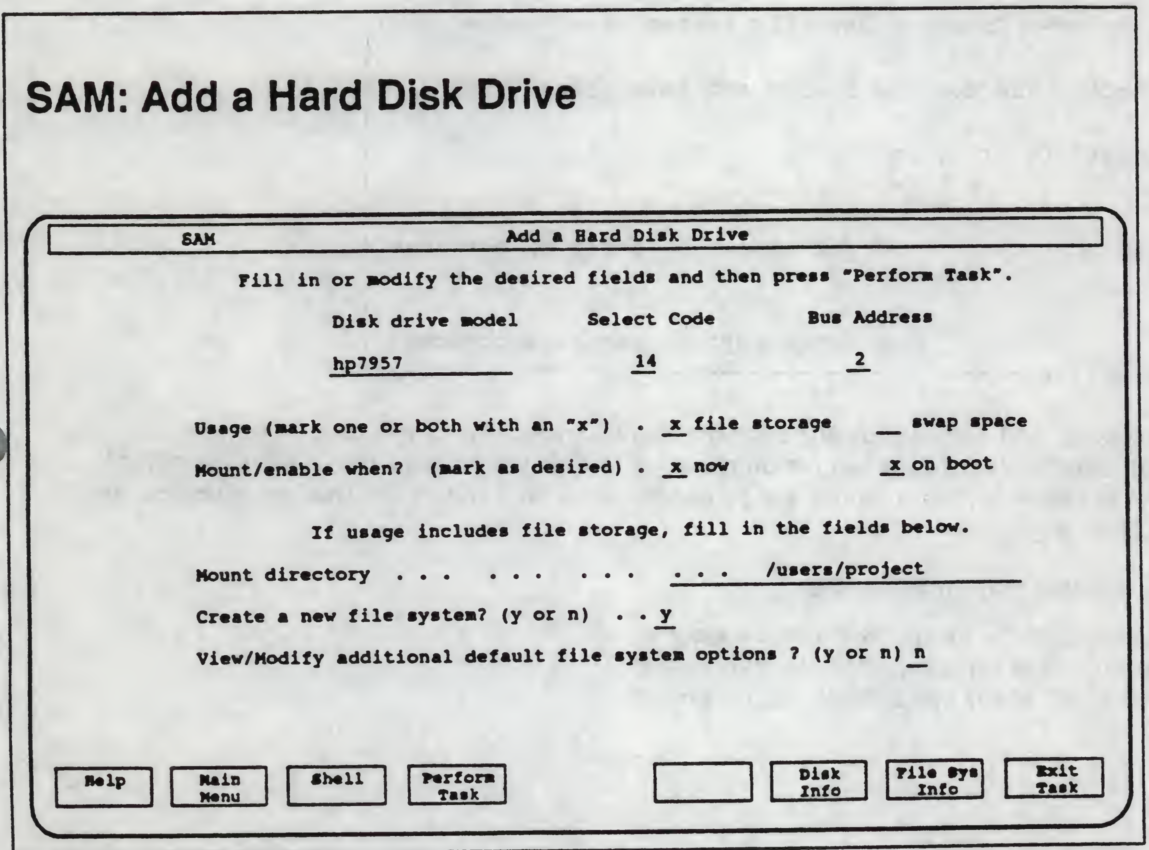


51436 CF-14. 14 

## Student Notes

You will see this screen if you selected "Add a Hard Disk Drive" option in the "Disk and Swap Configuration" menu. This allows you to add a new disk to your configuration and optionally create a new file system.

Before you see the menu in the slide there will be a pop-up window discussed in the last slide's Student Notes.

At this point SAM would have already filled in some of the particulars for the disk you are adding:

Disk drive model to be used as a label in `/etc/disktab`

Hardware addressing information

If you get confused about what is being asked for, SAM will "help" you if you place the cursor over a field in the menu and press the `Help` softkey.

# Module CF — File System Creation (Series 300/400/700)

During the SAM's process of adding a hard disk drive, various pop-up menus will appear asking about such things as whether to create a directory for a mount point that doesn't yet exist, and other options for `/etc/checklist`. One menu of particular interest is the following:

```
+-------------------- Create a New File System ----------------+
|                                                               |
| Fill in or modify the desired fields and then press "Done".  |
|                                                               |
| Initialize disk? (y or n)  n                                  |
|                            -                                  |
|                                                               |
| Disk space allocation:      M for swap,      M for file system |
|                         ---             ---                    |
|                                                               |
|                          (use arrow keys to see more choices) |
+---------------------------------------------------------------+
```

The initial values for swap and files system sizes are not shown while SAM is working on some calculations based on data tn `/etc/disktab` for on the type of disk you have. After a brief "working" message appears and is removed, some values are provided, along with indicators that more choices are available with the arrow keys.

Some of the choices for this particular drive are:

Disk space allocation: 0 M for swap, 78 M for file system
Disk space allocation: 12 M for swap, 66 M for file system
Disk space allocation: 20 M for swap, 58 M for file system

## CF-15. SLIDE: SAM: Add a Local File System

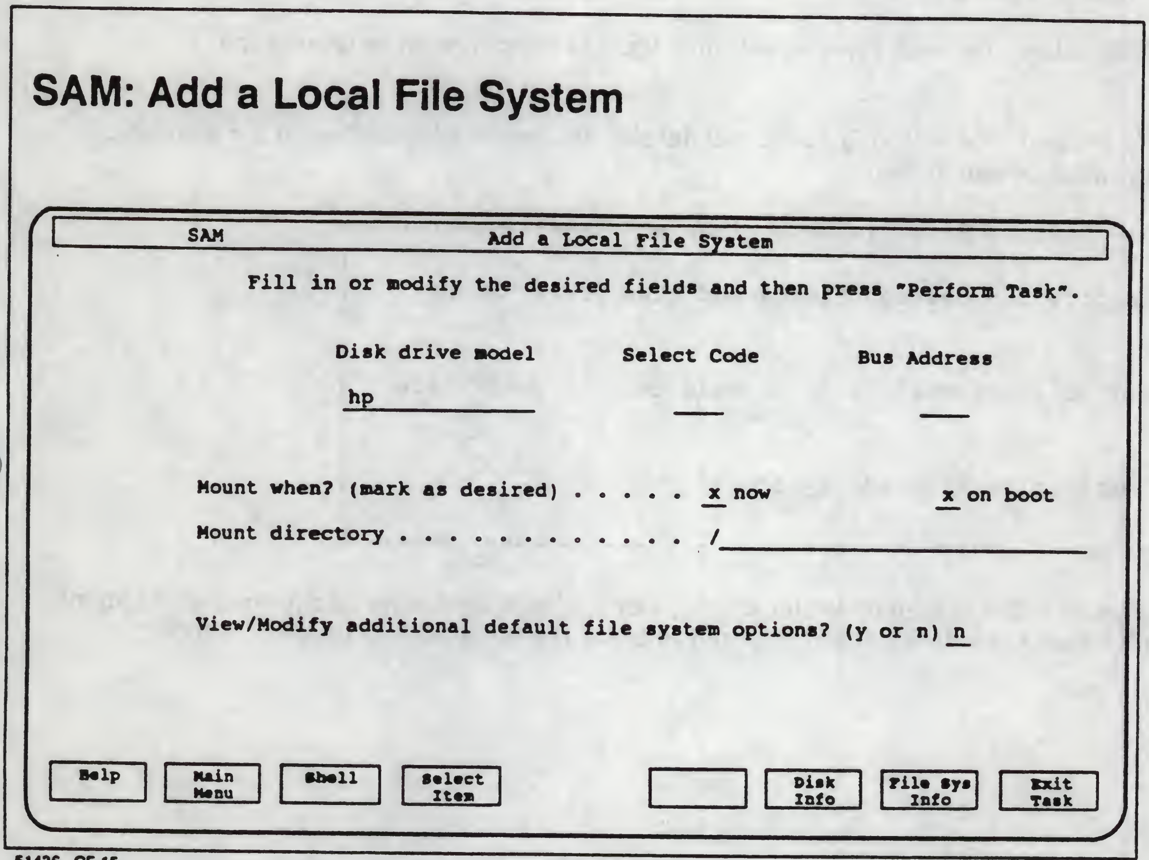


51436 CF-15. 15 

## Student Notes

You will see this screen if you selected "Add a Local File System" option in the "File Systems" menu. This allows you to add a file system to your configuration but does not have the capability to support creating a file system.

Before you see the menu in the slide there will be a pop-up window listing all the candidate file systems and disks that SAM could locate, along with their hardware addresses. The pop-up window will overlay the menu seen on the slide, you should select one of the choices by moving the arrow keys and pressing **Done**. If your disk is not shown, You can select **Exit Window**, get back to the menu in the slide, where you can exit to a shell.

If your disk is not displayed in the list, you should follow precautions described in the installation module to determine that your disk is correctly connected and configured. You may then select one of SAM candidates or fill in the drive model and hardware addressing information yourself.



147

Among the choices that SAM will list in the "Help" screen for the Disk drive model are: CDROM-HPIB and CDROM-SCSI. Allowing you to choose to add these to your file system is a key difference to the "Add a Hard Disk Drive" path we saw above.

Help doesn't list all the aliases for each type, so you may want to select a shell to browse the **/etc/disktab** file.

If you fill in the field marked "View/Modify additional default file system options?" with a y response, the following pop-up window will appear:

```
+-------------- Default File System Options --------------------+
|                                                                |
| Fill in or modify the desired fields and then press "Done".   |
|                                                                |
|                                                                |
| Write protection (mark one)  . ..    read only  x read/write  |
|                                      -          -              |
|                                                                |
| Set user ID execution allowed? (y or n)  y                     |
|                                                                |
+----------------------------------------------------------------+
```

This allows you to mount a disk read-only if you choose, along with setting other mount command options at mount time. Notice here there is less input required than for the other path in previous slides.

## CF-16. SLIDE: Disk Quotas

97

### Disk Quotas

- Limit the number of files a user can create.
- Limit the number of disk blocks a user can own.
- Established on a per-file-system basis
- Require superuser to set up and manage
- Enforce **soft limits** and **hard limits** for each user.
- Keep statistics on limits and usage by each user

51436 CF-16. 16 

## Student Notes

Disk **quotas** provide an improved means of controlling file system usage. This feature (new at HP-UX 8.0 release) sets user limits on the number of files created and disk blocks used on both NFS and HFS file systems. Quotas can enforce a **soft limit** and a **hard limit** for each user's file system usage. Disk quotas can be set up as limits for each user individually or for all users. Only a superuser can set up and manage disk quotas.

Users who exceed their soft limits are notified and have a certain amount of time to correct the problem. The time the user is given to correct the problem is also configured by the system administrator. Users who exceed their hard limits are or exceed their soft limits beyond the specified time interval are prohibited from creating files or increasing file system block usage. The following messages appears if a user reached the block limit:

```
DISK LIMIT REACED -- WRITE FAILED
```

The following messages appears if a user reached the file limit:



148

```
FILE LIMIT REACED -- CREATE FAILED
```

A system administrator must intervene to increase or decrease the set limits or the user may temporarily save his data to file system, `/tmp` for example, where he has not reached quota limits.

Statistics on user file system limits and usage are kept by the HP-UX operating system using data structures that reside in memory rather than on disk. This implementation reduces the overhead, and minimizes the "cost" of using quotas to an almost unnoticeable level.

## Pros and Cons of Disk Quotas

Advantages:

- When using disk quotas you control the way the file system is used.
- Impact on file system performance is rarely a problem.
- You can set a **soft limit** and a **hard limit** for each individual user, or groups of users, or for all users.
- You can set a **time limit** for each file system that uses quotas.
- You can set quotas running on any HFS or NFS file system. Not every file system needs to have quotas enabled.
- A user can check their quota status at any time.
- They are optional, you do not have to use quotas if you do not need to use them.

Disadvantages:

- When using disk quotas the reboot time of a crashed system may take longer.
- A system administrator must intervene to increase or decrease the set limits for any user.
- A system administrator must reset quotas to 0 before removing a userid from the system.
- You should not enable quotas on certain heavily used mounted file systems (eg `/tmp` directory)

The system administrator must add **dskquota** to the **dfile** in order to advantage of the functionality of disk quotas.

## CF-17. SLIDE:Readying the File System For Disk Quotas

9D

### Readying the File System For Disk Quotas

1. Mount the file system
2. Edit `/etc/checklist` to enable quota checking file
3. Create an empty file `quotas` in root level of the file system using `cpset(1m)`

of via opties achter mount .

51436 CF-17. 17 

## Student Notes

Disk quotas are set by executing the following steps:

1. Mount the file system.

2. Edit **/etc/checklist** to enable quota checking file. If the current entry exists it must be edited, otherwise a new entry must be created. To enable quota checking on a file system, this entry would need to be changed to:

**/dev/dsk/1s0 /users hfs rw,suid,quota 0 1**

The next time the system boots, disk quotas will be automatically enabled for that file system. Since the file system was mounted when this change occurred, disk quota checking is **not** currently enabled. You must reboot the system for it to take effect, or you may manually start up quotas as described later.



149

3. Create an empty file quotas in root level of the file system using cpset(1m).

```
# cpset /dev/null /users/quotas 600 root bin
```

The /users part of the file name specifies that the file is to be created at the top level or "root level" of the file system that will be mounted to the /users directory at boot time. The /dev/null specifies that the file is created empty. For further information on the cpset command, see the the cpset(1m) manual page in the *HP-UX Reference Manual*

This file will contain the limits and usage statistics after quotas is enabled on this file system. The data will be in binary form and keyed to each users "uid" number. The model for rapid access to this file uses the numerical value of the userid number to calculate the size of the data file. For this reason, you are encouraged to keep userid numbers small if you intend to enable disk quotas for that user. Fortunately, starting at the 8.0 release of HP-UX, SAM allows you to specify or change the userid assigned to new users on your system.

## CF-18. SLIDE: edquota: The Quota Editor

9-9

edquota: **The Quota Editor**

Invoke edquota for user block and inode to:

- Set up a "prototype limits" for a "prototypical user"
  /etc/edquota *proto_user*

  Sample input to edquota for user block and inode :
  fs /users blocks (soft = 1000, hard = 1200)
  inodes (soft = 250, hard = 300)

- Apply the prototypical limits to actual users
  /etc/edquota -p *proto_user username . . .*

Invoke /etc/edquota -t to set time limits fo˙ a file system

Example entry:

```
fs /user blocks time limit = 10.00 days, files time limit = 20.00 days
```

51436 CF-18. 18 

## Student Notes

**edquota** is the quota editor. The **edquota** command converts the binary data stored in the **quotas** file to a readable form so that the superuser can establish quota limits for users of the file system. **vi**, the editor of choice, is the default.

### User Block and Inode Quotas

Since the data required for quotas is arduous to type, often times a system administrator will set up a **prototype user** and merely reassign those limits to other users. What follows is an example on how to set up such a user. This is also the same syntax to assign limits to an actual user.

```
# /etc/edquota proto-man|
```

While inside the vi editor, you should type in a specification of the limits you desire. An example might be:



150

```
fs /users blocks (soft = 1000, hard = 1200) inodes (soft = 250, hard = 300)
```

Then you should write (save) the text file, and quit the editor. This entry sets block and inode limits for the file system **/users**.

If you want uniform limits for groups of users or for all users, apply the prototype limits to actual users. This is done with the **edquota -p** command. Given the example above we might type:

```
# /etc/edquota -p proto-man kathy ivan gisela robert
```

This command assigns the limits of the prototypical user **proto-man** to the users **kathy, ivan, gisela**, and **robert** who were already defined on the system as valid users. Notice that you can have several user names on the command line. You may also have several prototypes defined, but you use only one at a time with the **edquota -p** command.

---

**Note** If you plan to remove a user from the system, you must run **/etc/edquota** to set that user's limits to **0** before actually removing the user from the system.

---

### Time Limit Quotas

To set **time limits** for users, use the **edquota -t** command. Time limits allow the users a specific amount of time to reduce their usage when they exceed the soft limits. Time limits apply file-system-wide an to *all* users in a given file system.

You can set one time limit for correcting a violation of the soft limit on the number of files created and a different time limit on correcting violation of the soft limit on the number of blocks used.

To set the time limits type:

```
# /etc/edquota -t
```

While inside the **vi** editor, you should type in a specification of the time limits you desire. An example might be:

```
fs /users blocks time limit = 10.00 days, inodes time limit = 15.00 days
```

Then you should write (save) the text file, and quit the editor.

If you do not set these limits, the quota system defaults to seven days. If a user has exceeded his soft limit for more than the period specified by the time limit, the soft limit becomes the hard limit.

## CF-19. SLIDE: Turning Disk Quotas On and Off

100

### Turning Disk Quotas On and Off

Turn off quota checking:

/etc/quotaoff [-v] [-a|*filesys*]

Bring the quota file up to date:

/etc/quotacheck [-v] [-p] [-P] [-a|*filesys*]

Turn on quota checking:

/etc/quotaon [-v] [-a|*filesys*]

51436 CF-19. 19 

## Student Notes

**/etc/quotaoff** [-v] [-a|*filesys*] turns off quota checking for one or all file systems and optionally displays a message. Remember that you will have to run the **quotacheck** command to check disk consistency before you turn quota checking on again. **quotacheck** examines each specified file system and builds a disk-usage table of user IDs and blocks used, with it compares to a table stored in that file system's **quotas** file. If inconsistencies are detected, both the **quotas** file and the kernel quotas table are updated. Inconsistencies typically develop if a system mounted with disk quotas has been used for a period of time with quotas enforcement suspended. Since **quotacheck** accesses the raw device in calculating the actual disk usage for each user, the file systems checked should be inactive when **quotacheck** is executed. This is most easily done in single-user mode.

**quotacheck** recognizes the following options:

| | |
|---|---|
| **-a** | turn on quota checking for all file systems with **rw** and **quota** options in **/etc/checklist** |
| **-v** | report calculated disk quotas for each user on a particular file system |



151

-p check file systems in parallel as allowed by equal values in the *pass_number* field in /etc/checklist

-P Preen file system, checking only those files systems with invalid quota statistics. Also checks in parallel as -p option.

Example:

```
# /etc/quotacheck /dev/dsk/1s0
```

/etc/quotaon [-v] [-a|*filesys*] safely turns on disk quotas for a one or all file systems and optionally displays a message. As an example:

```
# /etc/quotaon -v /users
```

## CF-20. SLIDE: Quota Status

101

### Quota Status

```
# /usr/etc/quot /dev/rdsk/2s0

/dev/rdsk/2s0 (/users):
 2435    binky
  689    daffy
  345    bugs
```

```
# /usr/etc/quot -v /dev/rdsk/2s0

1793   bmc       1765   1272   1073
 296   vicki      295    295    222
  20   sandy        0      0      0
```

```
# repquota /users
                      Block limits                           File limits
User             used   soft   hard    timeleft   used   soft   hard       timeleft
bill      --       59    100    200                 24     30     40
fred      +-      199    100    200   1.7 weeks     10     30     40
joe       --       63    100    200                  9     30     40
dan       ++      173    100    200   1.4 weeks     32     30     40      1.4 weeks
```

51436 CF-20. 29 

## Student Notes

A user can get status report of his files by using quot or a system administrator can use quot to obtain information about all users on a file system by simply typing:

```
$ /usr/etc/quot /dev/rdsk/2s0

/dev/rdsk/2s0 (/users):
 2435    binky
  689    daffy
  345    bugs
```

See the man pages for more options.

However, a more interesting report can be obtained using `repquota` which provides the system administrator with a full status report.



152

## MF-1. SLIDE: File System Maintenance

1.02



51436/51482 MF-1. 145 

## Student Notes

One of the principal responsibilities of a system administrator is preservation of the user's data. Since the data storage structure utilized by HP-UX is the file system, it's imperative that the storage environment of the file system be checked regularly for possible problems. The integrity of the file system can never be compromised. There are many things the administrator can do to help maintain the integrity of the file systems, including:

- Use `fsck` regularly
- Ensure `sync` executes regularly
- Monitor disk usage
- Employ regular backup procedures
- Properly shut the system down

## MF-2. SLIDE: The BUFFER CACHE

103

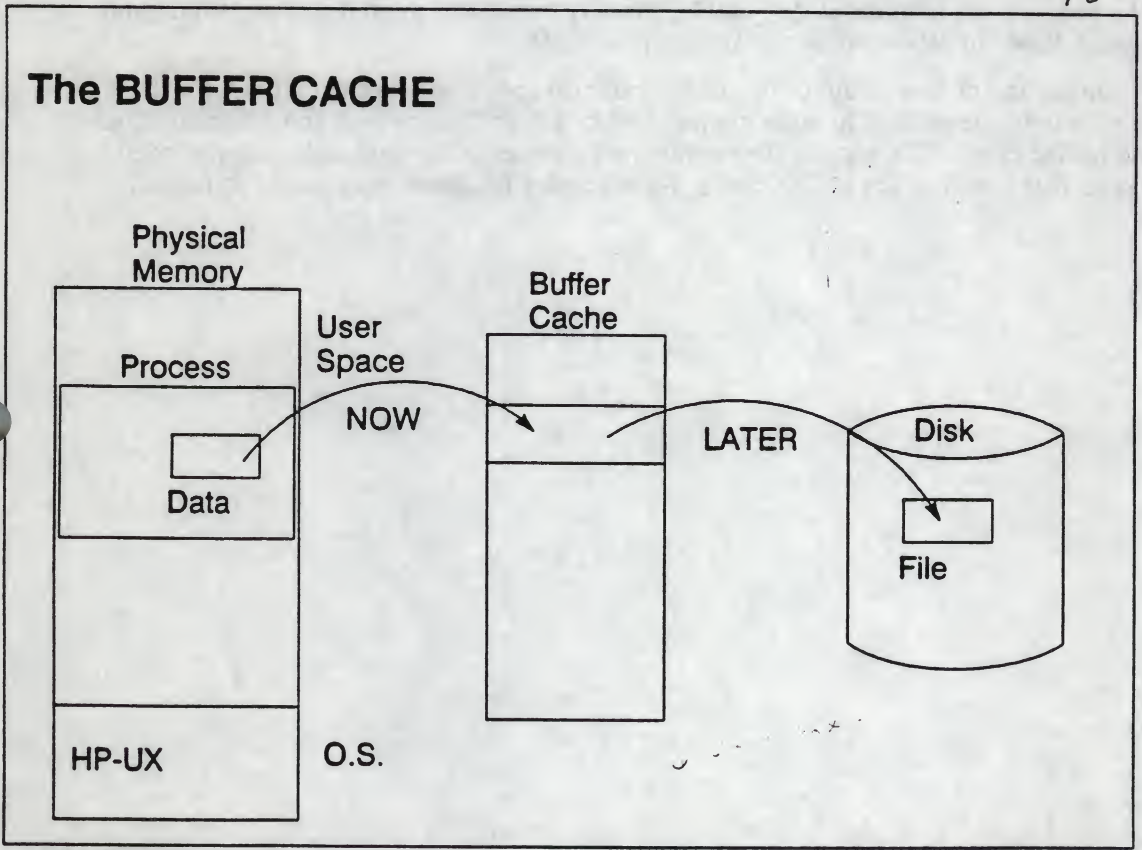


51436/51482 MF-2. 14C 

### Student Notes

When data needs to be written to a disk, the actual write does not occur immediately. The data is initially copied to an in-memory buffer called the **buffer cache**. This is a much faster operation than if the system had to perform an actual write to the disk. The data, along with the inode information, is written to the disk sometime later, usually when the buffer cache fills up and the system needs to clear some buffer space. If the system is halted without writing the buffer to disk, the file system could become corrupted. You should never allow work to continue if you suspect that the file system is corrupted.

Why does HP-UX bother with a **buffer cache**?? Well there are many benefits and advantages of having a **buffer cache**, so let us look at a few.

- The use of the buffer cache allows for uniform disk access, because the kernel does not need to know the reason for the I/O. The kernel just always writes buffers to the disk, not parts of buffers or real numbers. So system design is simpler, from a disk I/O standpoint.



155

- By using a buffer scheme programs are more easily ported to other UN*X systems. Disk I/O may be different on different UN*X machines, but the programs don't haveto know that, they simply write to buffers, without having to worry about how the disc is set up.
- Using a buffer cache reduces the amount of disc traffic, thereby increasing overall system throughput and decreasing reponse time. In other words the system runs faster.
- Re-use of data files and program files in the buffer cache can also speed up a system. (i.e. `vi` is a program file. It gets loaded in from disc to main memory when a user first invokes the program. There is also a copy in the buffer cache. The second time someone executes this program, it may not need to be loaded in from disc, if it is still in the buffer cache. So executing programs may be much faster.)

## MF-3. SLIDE: How HFS Handles File System Updates

104

### How HFS Handles File System Updates

- File data is first written to a buffer cache which:
  - Reduces amount of I/O necessary
  - Maximizes performance and efficiency
- Buffer contents are written to the file system when:
  - The buffer becomes full
  - The sync command is invoked
  - The file system is unmounted
- If the system is halted before buffers are written, the file system may become corrupted

51436/51482 MF-3. 147 

## Student Notes

Every time a file is modified, the HP-UX operating system performs a series of file system updates. For example, if data is added to a file which increases its size, additional block(s) may need to be allocated for the new data. The block(s) must be deleted from the free block list, the inode of the file must be updated, the superblock must be updated to reflect the new size of the file system, and countless other updates must occur. Now consider that this type of update sequence must occur for every change in the file system.

Updates occur to the superblock, inodes, data blocks, and cylinder group information in the following ways:

Primary Superblock — The superblock of a mounted file system is written to the disk whenever a `umount` command is issued, or when a `sync` command is issued and the file system has been modified. The root file system is mounted during boot and cannot be unmounted.

Inodes — An inode contains information specific to the file it describes. An inode is written to the file system upon closure of the file associated with the inode, when a `sync` or `fsync`



156

command is issued, when the file system is unmounted, or as soon as the file is written if `O_SYNCIO` is set for the file.

**Data blocks** In-core blocks are written to the file system whenever they have been modified and released by the operating system. More precisely, they are buffered or queued for eventual writing. Physical I/O is deferred until the buffer is needed by HP-UX, a `sync` command is issued, an `fsync` is issued for the file, or `O_SYNCIO` is set for the file. If a file is opened with the `O_SYNCIO` flag set, the `write` system call does not return until completed.

**Cylinder group information** The cylinder group information is updated whenever a `sync` is executed, or when the system needs a buffer and the cylinder group is written.

## MF-4. SLIDE: The sync Command

105

### The sync Command

- Writes buffer contents to disk
- Keeps the file system current
- Is normally invoked on a regular basis by the syncer program
  - `syncer` is executed in the `/etc/rc` file
  - The syntax of the `syncer` program is:

```
syncer [ seconds ]
```

51436/51482 MF-4. 148 

## Student Notes

As we have seen, data is written to an in-core buffer cache before it is written to disk. We have also seen that a physical write from the buffer to disk is delayed until:

- The system needs the buffer for another operation.
- The file system is unmounted.
- The sync command is executed.

When the sync command is executed, it causes the system to flush its buffers and write all data to disk.



157

## MF-5. SLIDE: Causes of File System Corruption

106

### Causes of File System Corruption

- Improper shutdown procedures
  - Not using `reboot` or `shutdown` to bring down the system
  - Taking a mounted file system off-line
- Improper startup procedures
  - Not checking a file system for inconsistencies
  - Not repairing inconsistencies found
- Hardware failure

51436/51482 MF-5. 149 

## Student Notes

A file system can become corrupt in a number of ways. The most common ways are improper shutdown procedures and hardware failures.

Although the hardware of a computer system is usually very reliable, the components of a computer system are very complex and fragile. Consequently, any piece of hardware could fail at any time. Hardware failures vary and can be as subtle as a bad block on a disk pack or as obvious as a non-functional disk controller. By following recommended hardware preventive maintenance procedures and by keeping regular backups, you can avoid most serious problems and be prepared for any that might occur.

## MF-6. SLIDE: The fsck Command

107

### The fsck Command

- Checks file system consistency and makes repairs
- Is multi-pass
- Should be run on an unmounted (or quiet) file system
- Can be run interactively or non-interactively
- Syntax:

```
fsck -p|-P [file_system]
```

autom.

```
fsck [-b block#] [-y | -n] [-q] [file_system]
```

51436/51482 MF-6. 150 

## Student Notes

**fsck** is the principal file system maintenance tool available with HP-UX. It verifies the structural integrity by checking data which is intrinsically redundant in a file system. The redundant data is either read from the file system or computed from known values.

**fsck** is a multi-pass program, meaning that it examines the file system a number of times, each iteration examining a different feature of the file system. Each pass **fsck** makes through the file system is known as a **phase**.

During each phase, any inconsistencies noted in the file system are reported and **fsck** asks if corrective action should be taken. **fsck** then waits for a response. If a yes response is provided, **fsck** attempts to repair the inconsistency. With a no response, **fsck** ignores the inconsistency found and continues its checking. There are very few occasions when a no response should be given.

**fsck** should be run on quiescent file systems. Ideally, the file system should be unmounted. Since this is not possible for the root file system, you should bring your system to a single-user run-level with the **shutdown** command before running **fsck** on root.

If you invoke **fsck** without giving it a *file_system*, **fsck** will check those file systems marked "hfs" in the **/etc/checklist** file. The order is determined by the number in the *pass_number* field (sixth field). The following are the two methods for invoking **fsck**:

/etc/fsck -p|-P [ *file system* . . . ]

/etc/fsck [ -b *block#* ] [ -y | -n] [ -q ] [ *file system* . . . ]

We will discuss both forms of the command.

**Note**

For preventive maintenance, it is a good idea to run **fsck** on all file systems once a week (or after each full backup).

## MF-7. SLIDE: fsck in Preening Mode

100

### fsck in Preening Mode

```
fsck -p [ file_system ]
fsck -P [ file_system ] ~~[illegible]~~ alles
```

Non-interactive mode

Fixes inconsistencies but does not remove data

Will correct:

- Unreferenced inodes
- Unreferenced pipes and fifos
- Incorrect link counts
- Missing blocks in the free list
- Blocks in the free list also in files
- Incorrect counts in the super-block

51436/51482 MF-7. 151 

## Student Notes

The first of the two ways that **fsck** can be invoked is preening mode. This is a non-interactive mode. This option fixes many problems, but never removes data. **fsck** decides what to do to fix a problem. For each problem it fixes, it prints a message identifying the file system and the corrective action taken. If it cannot solve a problem, it terminates. If **fsck** terminates, you should run it interactively so that you can fix the problems.

If the -p (lower-case) option is used without specifying a *file_system*, **fsck** reads the specified pass numbers in **/etc/checklist** to inspect groups of disks in parallel, taking maximum advantage of I/O overlap to preen the file systems as quickly as possible.

The -P (upper-case) option operates in the same manner as the -p option except those file systems which were cleanly unmounted will not be checked . (**fsck** checks the clean byte.) This can greatly decrease the amount of time required to reboot a system which was brought down cleanly. **fsck -P** is usually the command which is run at bootup time. It is invoked in the /etc/bcheckrc script run by init.

**Note**

It is a good idea to run `fsck -p` (lower-case) before doing a full backup. If `fsck -p` completes successfully, perform your backup as normal. If `fsck -p` aborts with errors, backup the bad file system, then attempt to repair the file system. If the repair completes successfully, backup again.

## MF-8. SLIDE: Other Options to fsck

109

### Other Options to fsck

fsck [ -b *block#* ] [ -y | -n ] [ -q ] [ *file_system* ]

-b Uses the block specified immediately after the flag as the super block for the file system

-y Assume a "yes" response to all questions asked by fsck

-n Assume a "no" response to all questions asked by fsck

-q Quiet mode

51436/51482 MF-8. 152 

## Student Notes

If you invoke fsck with no options, it runs interactively. In interactive mode, fsck poses a question when an inconsistency is found and waits for your response.

Other options available to fsck are shown on the slide. They are described in more detail below:

-y This option causes fsck to answer "yes" to all questions posed by the command. It is possible that data will be removed as a result of a "yes" answer. Consequently, if this option is to used, the file system should be examined with the -n option first so the possible consequences can be assessed.

-n This option causes fsck to answer "no" to all questions posed by the command. Since this option never results in the loss of data, it may appear safe. However, since "no" is supplied as the answer to all questions, fsck takes no corrective actions and inconsistencies are not resolved. It is recommended that this option be used when you want to assess the state of the file system; however, you should invoke the command again to resolve inconsistencies if any are found.



160

Example:

```
# fsck -n /dev/dsk/c0d0s3  |  tee /usr/tmp/fsck.log
```

In this example, the diagnostic output is directed to the file `/usr/tmp/fsck.log` as well as to the screen. This output can be analyzed to determine corrective action.

`-q` This option causes `fsck` to print only those messages that require a response.

`-b` *block#* This option tells `fsck` to use *block#* as the superblock for the file system check. This is useful if the primary superblock is lost or corrupted. You could try this option if you seem to be getting a lot of errors with `fsck` and it doesn't make sense to you.

You can do a Ctrl-c out of `fsck`. `fsck` does not make any repairs to a file system until the command successfully completes. (Repairs are written to RAM as the command is running.) Thus, if you kill the command, nothing will happen.

**Note** Do not reboot the system unless `fsck` tells you to reboot. A reboot will `sync` the disks and thus write out the bad data. If you must reboot, use the `reboot -n` command which does not issue a `sync`.

## MF-9. SLIDE: The lost + found Directory

11c

### The lost+found Directory

- Created by newfs (it should exist in every file system).
- Must have some empty file slots.
- fsck copies problem or orphaned files to this directory.
- Check this directory after each invocation of fsck and try to determine the origin of any files there.
- Can be created with the mklost+found command. (mklost+found creates empty slots.)

51436/51482 MF-9. 153 

## Student Notes

Every file system should have a lost+found directory at the root of its file system (that is, /lost+found). The lost+found directory is created by newfs. However, you should verify that the directory exists before using fsck to check the file system. If lost+found does not exist, you can rebuild it with the /etc/mklost+found command.

The lost+found directory must have some empty slots so that fsck can write to it. This is best accomplished by copying a number of files to lost+found and then removing them. For example:

```
# cd lost+found
# touch `ls /bin`
# rm *
```

This only needs to be done before fsck is invoked for the first time.



161

**fsck** places any problem files or directories in the **lost+found** directory. After **fsck** completes, you should examine the contents of the directory. The files that are placed there should be moved back to their original directories. It may be difficult (and sometimes impossible) to determine where the files actually belong, but you should try to find owners. Run the **file** command on a file. If the file contains text, look at its contents to try and determine the owner. If the file contains executable code check to see if it has an SCCS identification string. If it does, the **what** command will list SCCS identification information. If the file does not have an SCCS identification string, use the **strings** command to print the literal strings from the file. These strings may help identify the owner. Do not execute an executable file found in the **lost+found** directory to try and figure out what it is. It may be the program that corrupted the disk.

## MF-10. WORKSESSION: Review Questions

1. Describe the Buffer Cache.

2. What are the advantages of the Buffer Cache?

3. What one disadvantage of the Bufer Cache does the System Administrator have to worry about?

4. What does `fsck` do?

5. How can a system administrator prevent having a corrupt file system?

6. What is the `lost+found` directory used for?



162

## MF-11. LAB: Lab Exercises (Series 300 & 800)

### Directions

Work with your lab partner to perform the following exercises.

1. Look at the `/etc/bcheckrc` script to determine the mode of `fsck` to be run upon bootup.

Will the `-P` option to `fsck` look through a clean file system?

When and how is the `/etc/bcheckrc` script executed?

2. On which file systems would the command `fsck -p` perform a file system check? In what order would these file systems be checked? Why is the order potentially significant ?

3. Run `fsck` on the file system you created and mounted in the module on Creating Your File System.

# Module BU — Backing Up the System

## Objectives

Upon completion of this module, you will be able to:

- Describe several backup strategies for a system.
- Differentiate between different backup and restore methods.
- Backup and restore files with `fbackup/frecover`, `cpio`, `tar`, and `dd`.
- Find files with the `find` command.
- Use `tcio` to buffer I/O to and from a cartridge tape drive.
- Use SAM to perform file system backups



163

# Module BU — Backing Up the System

## BU-1. SLIDE: Why Backup?

### Why Backup?

How Much Data Can You Afford To Lose?

Data is sometimes lost by:

- File system corruption
- Accidental removal of files
- Hardware failures
- System crash

Regular backups:

- Minimize data loss
- Keep users happy
- Provide stability and order

51436/51482 BU-1. 1 

## Student Notes

As we mentioned near the beginning of the course, one of the principal responsibilities of a system administrator is preserving the data stored on the system. Unfortunately data is sometimes lost. A piece of hardware may fail, a file is accidentally removed or overwritten, a command may go astray, or the system may crash. The user community has a reasonable expectation that the administrator has planned and implemented regular backup procedures to minimize the data loss.

To further minimize the chance of data loss, all backup media should be stored at a location geographically distinct from the system's disk drives. Picture the scenario of a multi-year project under development and a system administrator who has dutifully maintained system backups since the project's inception. However, for ease of retrieval, the backup media are stored in the same room as the computer system. If a fire, flood or other disaster destroyed the computer room, all the work done on the project would be lost. Unfortunately, it happens. Consequently, the safest course of action is to store the backup media in a secure environment separate from the computer equipment. In some cases this may mean the next room, a data vault on the premises, or even a data vault at another site.

# Module BU — Backing Up the System

The exact backup procedure employed is determined by a number of factors. A heavily used system, both in terms of number of users and amount of activity, may require some form of backup to be conducted daily. A lightly used system may only need to be backed up weekly or bimonthly. Media used for backups may prove to be expensive. If complete system backups are performed daily, the costs in terms of materials and personnel may grow to be significant. Consequently, a weekly or bimonthly backup may result in minimal costs added to the maintenance of the system.

In this module, we'll look at different backup strategies. Though each has its advantages and disadvantages, some backup strategy should be implemented as soon as users begin to work on the system.



164

## BU-2. SLIDE: Backup Strategies

112

### Backup Strategies

- Backup the entire file system (full backup)
- Backup part of the file system
  - Files that have changed since the last backup (incremental or delta backups)
  - A subtree of the file system
- Use a mixed strategy
  - Full backup once a week
  - Incremental backups daily

51436/51482 BU-2. 2 

## Student Notes

Before we can discuss *how* to back up the system, we must first discuss *when* to make backups and *what* to back up. There are two lines of thought concerning what to back up. You can selectively back up parts of the system, or you can back up the entire system. Most of the time, you should implement a backup strategy based on a combination of these two strategies. As an example, you could do a full system backup (or archive backup) once a week, on Sunday, and do incremental (delta) backups once each day on Monday through Friday. During the incremental backup, you may choose to back up only those files that have changed since the last full backup.

There are three things to consider when choosing a backup strategy that is right for you:

- The amount of media that it will take
- The amount of time that it will take
- How often you must make backups

## BU-2. SLIDE: Backup Strategies

43

### Backup Strategies

- Backup the entire file system (full backup)
- Backup part of the file system
  - Files that have changed since the last backup (incremental or delta backups)
  - A subtree of the file system
- Use a mixed strategy
  - Full backup once a week
  - Incremental backups daily

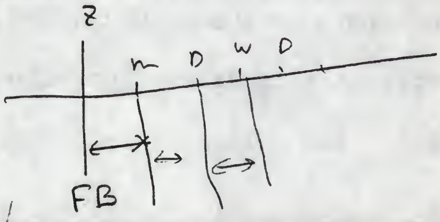


51436/51482 BU-2. 2 

beter; altijd van FB uitgaan → Δ groeit.

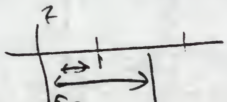


## Student Notes

Before we can discuss *how* to back up the system, we must first discuss *when* to make backups and *what* to back up. There are two lines of thought concerning what to back up. You can selectively back up parts of the system, or you can back up the entire system. Most of the time, you should implement a backup strategy based on a combination of these two strategies. As an example, you could do a full system backup (or archive backup) once a week, on Sunday, and do incremental (delta) backups once each day on Monday through Friday. During the incremental backup, you may choose to back up only those files that have changed since the last full backup.

There are three things to consider when choosing a backup strategy that is right for you:

- The amount of media that it will take
- The amount of time that it will take
- How often you must make backups



165

The amount of media and the amount of time require that you be able to approximate the amount of space being used by files on the system. You can use the bdf command to get the number of free disk blocks. The bdf command reports in 1024-byte blocks.

Series 300/400/700:

```
$ bdf
Filesystem    kbytes  used    avail   capacity  Mounted on
/dev/dsk/0s0  487022  142518  295801    33%     /
```

Series 600/800:

```
$ bdf
Filesystem         kbytes  used    avail   capacity  Mounted on
/dev/dsk/c0d0s13   580734  375387  147273     72%    /
/dev/dsk/c3d0s2    580318  173800  348486     33%    /users
```

How often that you make backups depends on how much data you can afford to lose. If you can afford to lose a month of data, then you need only back up the system once each month. If you can only afford to lose 6 hours of data, then you must back up every 6 hours. However, backing up every 6 hours can become prohibitive, and other possibilities (such as redundant systems) must be considered. For most applications, full backups once each week and partial backups each night are sufficient.

```
                         Example: Mixing Backup Strategies

  Sun      Mon      Tues     Wed     Thur     Fri      Sat    Sun      Mon

Archive     |         |        |        |        |        |   Archive   |
|           |         |        |        |        |        |      |      |
|<--------Incr        |        |        |        |        |      |<-- Incr
|<---------------   Incr       |        |        |        |
|<------------------------- Incr        |        |        |
|<---------------------------------- Incr        |        |
|<------------------------------------------- Incr        |
|<--------------------------------------------------- Incr
```

Paste figure BU-1 here. (Example: Mixing Backup Strategies)

Figure BU-1.

## BU-3. SLIDE: Comparison of Backup Methods

114

### Comparison of Backup Methods

universal

Table 1-1.

| | fbackup | cpio | tar | dd |
|---|---|---|---|---|
| Backup Type | Logical | Logical | Logical | Physical (image copy) |
| Speed | Fast | Medium | Slow | Fast |
| Flexibility | High | Medium | Low | Low |
| Handles Multiple Media | Yes | Yes | Yes (on HP-UX) | Yes |
| Portability | HP-UX only | Unix | Unix/some Non-Unix | N/A |
| Special files | Yes | Yes | No | Yes |
| Network "smart" | Yes | No | No | No |
| Can Append Files | No | No | Yes | Yes |
| Interruptible | Yes | No | No | No |
| Directory Relative Recovery | Yes | If created with relative paths | If created with relative paths | No |
| Application | Full and incremental system backups | Copying directory structures and transferring files to other systems | Easy to use and highly portable | Duplication of bit images on to different media |

51436/51482 BU-3. 3 

## Student Notes

All these utilities allow saving files to backup media. Below is a discussion of each:

**fbackup** Written specifically to be used in HP system backups, `fbackup` allows the use of 9 different backup levels, thus making the backup process extremely flexible. `fbackup` and its companion `frecover` allow an "across the network" back up. The commands are considered "Network smart", in that they allow you to specify a device on a remote host as part of the backup/recover command line. As noted on the slide this is an HP proprietary backup method usable on HP-UX systems only.

Another significant advantage of `fbackup` is that is now supported on all HP 9000 HP-UX systems as a feature of SAM. This means that series 300/400, series 700, and series 600/800 systems can use the same tools for easy backup and recovery of files.

**cpio** Being a general purpose Unix logical file backup utility, `cpio` is very common in the Unix world. This was the backup method of choice before `fbackup` because it also allows for backups to span multiple backup. `cpio` is option oriented and it has the capability for



166

selective file restores and is much faster than `tar` for large amounts of data. `cpio` also can backup special files such as device or network files.

At the HP-UX 8.0 release, `cpio` is POSIX compliant when you use certain options upon creating or reading archives. The default behavior is POSIX, an option allows backwards compatibility.

`tar` Also a general purpose logical file backup utility, `tar` is the oldest of all the utilities and thus the most portable. It is on all Unix systems and can be read by several non-Unix systems as well. `tar` is much easier to use than `cpio` and allows appending to backup media. Prior to HP-UX 7.0 release, `tar` did not permit the backup to span multiple media. This limitation may still be found on other vendors' system. Thus, `tar` has historically been used for backing up small numbers of files easily. `tar` also allows selective restore of files.

At the HP-UX 8.0 release, `tar` is POSIX compliant when you create or read tar archives. The default behavior is to create and read POSIX format archives of files. You must use a key, 0 to get the old behavior for backwards compatibility. You will notice other differences in `tar` behavior with the v verbose key.

### Another "Backup" Utility, dd

The dd command is useful in some limited situations, but is technically *not* a backup command. The dd command is a general purpose physical file copy utility. This is different from all the above utilities in that dd copies no file names or file attributes to a backup media; it simply copies everything, bit for bit. Thus no selective restore is possible. For these reasons, dd is *not* recommended for system backups. It is generally used for two purposes:

- Make a duplicate copy of a disk quickly. This assumes a destination disk the same size or larger than the source.
- Read or translate a foreign 9-track magnetic tape. For example, dd has the capability for reading backup media with user defined record sizes, ASCII or EBCDIC translation, byte switching, and other useful options.

---

**Note** Since there are many commands available, and many options to each of these commands, it is strongly recommended that you write the command you use to create your backup on the label of your backup media.

---

## BU-4. SLIDE: Alternative Backup Methods

115

### Alternative Backup Methods

Voor method.

Table 1-2.

| | ftio | dump | rdump | backup | pax |
|---|---|---|---|---|---|
| Backup Type | Logical | Logical | Logical | Logical | Logical |
| Speed | Very fast | Fast | Fast | Fast | Fast |
| Flexibility | High | Medium | Medium | Medium | Medium |
| Handles Multiple Media | Yes | Yes | Yes | Yes | Yes |
| Portability | cpio compatible | BSD Unix | BSD Unix | uses cpio | POSIX compliant, cpio and tar compliant |
| Special files | Yes | Yes | Yes | Yes | Yes |
| Network "smart" | Yes | No | Yes | No | No |
| Can Append Files | No | No | No | No | Yes |
| Interruptible | No | Yes | Yes | No | No |
| Directory Relative Recovery | Yes | Yes (depends) | Yes (depends) | Yes (depends) | Yes |
| Application | Full and incremental system backups | Full and incremental system backups | Full and incremental networked system backups | Full and incremental system backups | POSIX data interchange |

51436 /51482 BU-4. 4 

## Student Notes

All these utilities allow saving files to backup media. Below is a discussion of each:

**ftio** The "Fast Tape IO" program was written specifically to be used in HP system backups as a faster alternative to `cpio` backups. It is not suggested as a primary system backup method now that `fbackup` is available. `fbackup` is more flexible and also a supported feature of SAM. The `ftio` command allows the use of 9 different backup levels, thus making the backup process more flexible then `cpio`, for instance. The `ftio` command has been designed to be compatible with the `cpio` so that you can read or write compatible backup media with this utility.

**dump** This program is specifically supported to be used where there are other computers with a Berkeley (BSD) based Unix, and expertise already exists using the Berkeley dump and `restore` command. The dump utility allows the use of 9 different backup levels, thus making the backup process very flexible.



167

| | |
|---|---|
| `rdump` | The dump utility and its companion `restore` have been enhanced to allow an "across the network" back up. These commands are called `rbackup` and `rrestore`. The leading "r" stands for remote. `rbackup` is considered "network smart" since it allows you to specify a device on a remote host as part of the command. This implies a networked system where a cooperating process can run on a separate host. This utility is to be used where the customer has other computers with a Berkeley (BSD) based Unix, and is accustomed to using the Berkeley `rdump` and `rrestore` command. The `rdump` utility also allows the use of 9 different backup levels, thus making the backup process very flexible. |
| `backup` | This is a shell program or script designed for ease of use. It is a general purpose file system backup tool provided by HP. The script uses the `cpio` command, and can be used "as shipped", or modified by you for customized application. This was the backup method of choice before `fbackup` because it handles full archive and incremental backups with the same command. There are only a few options, so it is simple to use. It is well documented with comments, which makes it easy to customize. |
| `pax` | Portable Archive eXchange (PAX) is a utility which is emerging as a future industry standard general purpose backup utility. The command `pax` is the newest of all the above utilities. It was specified by a US government computer standards committee as POSIX.2 Draft 9 in the interest of developing a computer industry standard backup and file interchange mechanism. Since the standard has not been finalized, `pax` is *not* widely used for system backups, it may become popular over time. `pax` reads and writes data in the *POSIX Archive/Interchange File Format*, and also supports Version 7 `tar` and System V `cpio` formats. HP does not recommend using this implementation of `pax` as the primary backup mechanism since there may be some changes in the standard as it evolves. |

## BU-5. SLIDE: The find Command

116

### The find Command

Syntax:

find *path-list expressions*

| | |
|---|---|
| *path-list* | the list of directories to search recursively |
| *expressions* | search criteria and actions |

Examples:

```
find / -print
find /users -name .profile -print
find . -size +2000 -print
find / -atime +30 -print
```

± 1 meg.

51436 /51482 BU-5. 5 

## Student Notes

Before discussing any of the backup commands in detail, we will discuss a command that is frequently used in conjunction with backup commands, `find`.

The `find` command can be used to find a group of files that match a certain criteria. The `find` command is one of a very few commands that will perform an automated search through the file system. This ability makes it extremely versatile. Unfortunately, since the basic operation of `find` is to conduct a search, it is very slow and consumes considerable system resources.

The format of the command is:

`find` *path-list expression*

The *path-list* is a list of path names, typically one directory. Often "." is specified which results in a search starting in the current directory. A *path-list* of "/" results in a search of the entire file system. The path names are searched recursively for files which satisfy the criteria specified in *expression*. When



168

`find` locates a match, it performs the tasks also specified in *expression*. One of the most common tasks is to print the path name to the match.

The *expression* is made up of keywords and arguments which can specify search criteria and tasks to perform upon finding a match. One of the things which makes `find` so complicated is that the keywords used in *expression* are all preceded by a "-", so it looks like the arguments precede the options.

Examples:

`$ find / -print`
finds all files in the file system and prints their names on the screen

`$ find /users -name .profile -print`
finds all the files named `.profile` searching recursively down beginning at the directory `/users`

`$ find . -size +2000 -print`
finds all the files beginning at the current directory that are larger than one megabyte (2000 512-byte blocks).

`$ find / -atime +30 -print`
finds all the files in the entire file system that have not been accessed for 30 days

`find / -name core -exec rm {} \; -print`
finds all files named `core` (core dumps) in the entire file system and removes them. Notice the space following the curly braces, it is required.

The list of expressions and options for the find command can be found on `find(1)` man page in the *HP-UX Reference* manual.

## BU-6. SLIDE: The cpio Command

117

### The cpio Command

- Syntax:

```
cpio -o [ cvxB ]
cpio -i [ cvxdumB ] [ patterns ]
cpio -p [ vdumxl ] directory
```

- Reads a list of file names from standard input
- Copies contents of each file to standard output
- Examples:

```
# find / -print | cpio -ocx > /dev/rmt/0m  (backup)

# cpio -icxvmd "*/users/*" < /dev/rmt/0m   (restore)
```

51436/51482 BU-6. 6 

## Student Notes

The cpio command (with the -o option) reads a list of file names from standard input and copies the contents of the files specified to standard output. Thus, we must use cpio within a pipeline with some other HP-UX command generating the list of file names.

The cpio command is used in three major ways:

| | |
|---|---|
| -o | Read standard input and copy each file to standard output (make a backup) |
| -i | Read standard input for the backup data and recreate it on the disk (restore backup) |
| -p | Read standard input for file names and recreate those files in another directory (pass option) |

There are also several additional options that may be used with each major option:



169

**Table BU-1.**

| -o -i -p | Function |
|---|---|
| -c -c — | Write header in ASCII format (if used with -o, it must be used with -i) |
| -x -x -x | Handle special (device) files |
| — -u -u | Unconditionally restore (normally, an older file will not replace a newer file) |
| — -m -m | Retain current modification date (important for version control) |
| — -d -d | Recreate directory structure as needed |
| -v -v -v | Print a list of files copied to the terminal |
| — -t — | Print only a table of contents of the input |
| -B -B — | Block input/output at 5120 bytes to the record |
| -h — — | Follow symbolic links as if they were normal files |

Magnetic tape examples of the `cpio` command are as follows:

`# cd /users`
`# find . -print | cpio -ocx > /dev/rmt/0m` This example creates a relative backup on `/dev/rmt/0m`. Files can be restored to any disk directory.

`# find /users -print | cpio -ocx > /dev/rmt/0m` creates absolute backup on `/dev/rmt/0m`. Files can be restored only to the `/users` directory from where they were copied.

`# cpio -iudmcx < /dev/rmt/0m` restores backup to working directory if the tape was created in relative format.

`# find . -print | cpio -pdm /tmp` This example copies the subtree rooted at current working directory to subtree rooted at `/tmp`.

## BU-7. SLIDE: The `tcio` Command

118

### The tcio Command

- Syntax:

```
tcio -o [ dervVZSlnT ] file
tcio -i [ drvZSlnT ] file
tcio -u [ rvVml ] file
```

- Used with `cpio` when using cartridge tape
- Buffers I/O to relieve strain on tape drive
- Examples:

kostant altijd

```
# find . -print | cpio -ocx | tcio -o /dev/rct/update.src

# tcio -i /dev/rct/update.src | cpio -icdvmx

# tcio -i /dev/rct/c3d0s2 |  cpio -itcx

# tcio -uVr /dev/rct/c3d0s2
```

51436/51482 BU-7. 7 

## Student Notes

The wear and tear on the tape drive is extensive when redirection is used because the data transfer rates between the host computer and the cartridge tape drive are not synchronized. Thus, the `tcio` command was written to "buffer up" the data transfer between `cpio` and the cartridge tape drive. Instead of redirecting the output of `cpio` straight to the device, it is piped through `tcio` to enable this streaming to occur.

The `tcio` command, like `cpio`, has three major options:

| | |
|---|---|
| -o | go out to a device |
| -i | come in from a device |
| -u | utility option |

The `-o` and `-i` options of `tcio` correspond to those used with `cpio`. The `-u` option is for sending commands to the cartridge tape drive.



170

Examples of the tcio command are as follows:

- Make relative backup on cartridge tape /dev/rct/update.src
  `# find . -print | cpio -ocx | tcio -o /dev/rct/update.src`
- Restore relative backup to current working directory
  `# tcio -i /dev/rct/update.src | cpio -iudmcx`
- Extract a table of contents
  `# tcio -i /dev/rct/c3d0s2 | cpio -itcx`
- Unload tape from /dev/rct/c3d0s2
  `# tcio -uVr /dev/rct/c3d0s2`

## BU-8. SLIDE: The `tar` Command

119

### The `tar` Command

Syntax:

```
tar key [ args ] [ file | -C directory ]
```

Where *key* could be the following:

| | |
|---|---|
| `f` *dev-file* | Causes `tar` to use next argument as name of device where the archive will occur (this key is required) |
| `c` | Creates a new archive |
| `r` | Appends to end of archive |
| `x` | Extracts from archive |
| `t` | List table of contents of archive |
| `v` | Verbose, type the name of each file affected |

Examples:

```
# tar cvf /dev/rmt/0m file1 file2

# tar xvf /dev/rmt/0m

# tar cvf -  -C /users -C /etc | tcio -o/dev/update.src

# tar rvf /dev/rmt/0m file2
```

51436 /51482 BU-8. 8 

## Student Notes

`tar` saves and restores archives of files on a magnetic backup media (DDS format DAT, mag tape, or cartridge tape), optical disk (MO), flexible disk, or regular file. Its actions are controlled by the *key* argument. The *key* string can be preceded by a hyphen (-) (as when specifying options in other HP-UX commands), but it is not necessary.

If the `f` key is used, it causes `tar` to use the next argument as the name of the archive instead of `/dev/rmt/0m` (which is the default). If the name of the file is `-`, `tar` writes to the standard output or reads from the standard input, whichever is appropriate. Thus, `tar` can be used as the head or the tail of a pipeline.

The -C option can be included in the *file* list. It causes `tar` to change to the given directory and archive the files there. This allows multiple directories not related by a close or common parent to be archived using path names.



171

When the end of tape is reached, `tar` prompts the user for a new special file and continues. If a nine-track tape drive is used as the output device, it must be configured in Berkeley compatibility mode.

For a complete list of tar keys and options, see `tar(1)`.

The tar command has new default behavior. At the 8.0 release, `tar` is POSIX compliant when you create or read tar archives. The default behavior is to create and read POSIX format archives of files. You must use a key, 0 to get the old behavior for backwards compatibility. You will notice other differences in `tar` behavior with the `v` verbose key.

## BU-9. SLIDE: The fbackup Command

120

### The fbackup Command

- level f

Syntax:

`fbackup -f device [-0-9] [-u] [-i path] [-e path] [-g graph]`

| | |
|---|---|
| -f *device* | the device to which output will be sent |
| [ -0-9 ] | backup level — default is 0 |
| [ -u ] | causes fbackup to update the /usr/adm/fbackupfiles/dates file (only if used with -g) |
| [ -i *path* ] | causes path (file or directory) to be included in the backup |
| [ -e *path* ] | causes path (file or directory) to be excluded in the backup |
| [ -g *graph* ] | *graph* is a file that contains a list of files and directories to be included or excluded from the backup |
| [ -I *path* ] | write an index to file *path* |

+ tcio !!

Examples:

```
# fbackup -u8f /dev/rmt/0m -g /usr/adm/fbackupfiles/graphfile
# fbackup -f /dev/rmt/0m -i . -e ./subdir
```

51436 /51482 BU-9. 9 

fbackup -1f- | tcio -o /dev/update.src

## Student Notes

The **fbackup** command is the primary tool for creating both full and incremental backups. The **fbackup** command is quite flexible and allows each system administrator to develop a backup strategy that best suits the needs of the installation.

The general form of the command is shown on the slide. The first example on the slide backs up files and directories indicated in the **graphfile** to high speed reel tape. The second example backs up all files under the current directory except those contained under **subdir**.

The **fbackup** command does not, by default, write to the standard output and the **-f** *device* part of the command is not optional. The *device* may be a regular file or a device (special) file. You may specify - as the *device* to have **fbackup** write to the standard output. For most systems, a magnetic tape drive is used as the backup device. In this case, use the appropriate device special file for the raw (character) device in the */dev/rmt* directory.

## Backup Levels

The [ 0-9 ] option allows the user to define and use what are referred to as backup levels. Recall that an incremental backup strategy only makes backup copies of files that have changed. The key question to answer is "changed since *when*"? Certainly, any file on the system can be considered to have changed if the file's time stamp is compared to the beginning of time. In other words, an incremental backup that copies files that have been modified more recently than the beginning of time is actually a full backup!

Level 0 — is pre-defined to mean "the beginning of time" by the `fbackup` command. If invoked without a backup level option, `fbackup` performs a level 0 backup that results in a full backup.

Levels 1 - 9 — are used for incremental backups. Each time `fbackup` is invoked with a backup level option, `fbackup` makes copies of files that have changed since the last time a backup was made at a *lower* level. For example, suppose `fbackup` is run at level 0 on Monday, level 1 on Tuesday, and level 2 on Wednesday. A full backup would be made on Monday, an incremental backup of files that had changed since Monday would be made on Tuesday, and an incremental backup of files that had changed since Tuesday would be made on Wednesday.

It is up to the system administrator to assign meaning to each backup level. The backup strategy mentioned on the previous page called for full backups on Monday and incremental backups dating back to the most recent full backup on Tuesday, Wednesday, Thursday, and Friday. This scheme could be implemented by making a level 0 backup on Monday and level 5 backups on Tuesday, Wednesday, Thursday, and Friday. The reason for selecting level 5 for the incremental backups is to allow for some flexibility in the backup scheme should changes be required.

## Graph Files

HP-UX file system trees can get very large. Multiple physical disks can be mounted under a single file system to produce a logical file system tree that is enormous. For this reason, it is desirable to have a mechanism to specify only parts of the file system to be backed up. It may also be the case that part of a file system is entirely static and may require only occasional backup (for example, monthly or semiannually). Simply put, you may wish to specifically include or exclude parts of the file system during backup.

Inclusion and exclusion is accomplished either through the use of graph files (`-g` option) or through the inclusion of parameters on the command line (`-i` and `-e` options). A graph file is a file that contains ASCII text. Each line in the file contains a directory path that is to be included or excluded. Lines that begin with an "`i`" indicate a directory path that should be included. Lines that being with an "`e`" indicate files that will be excluded.

```
i /usr
e /usr/tmp
e /usr/lib
i /users
e /users/guest
```

The combination of backup levels and graph files provide significant flexibility. Recall that `fbackup` uses backup levels as a mechanism to identify files that have changed since the most recent backup that was made at a lower backup level. When invoked with a backup level *and* a graph file, `fbackup` makes a backup of files that changed since the most recent backup of that graph at a lower level.

### Other Common Options

As described above, fbackup is capable of determining when a backup of a specific level was last made. This implies that there must be a data file used by fbackup to retain the necessary information. Such a data file does exist and is readable ASCII text. The file is /usr/adm/fbackupfiles/dates. Since it may be desirable to create a backup without updating the dates file, fbackup requires the -u option to be specified on the command line if the user wants to update the dates file.

The [ -I *path* ] options causes fbackup(1M) to write an index to file *path*

There are other options that may be of interest to you. The manual page for fbackup(1M) provides explanations for all the options.

### Examples

```
# fbackup -u8f /dev/rmt/0m -g /usr/adm/fbackupfiles/graphfile
```

Perform a level 8 backup to /dev/rmt/0m using the graph file /usr/adm/fbackupfile/graphfile. Update the /usr/adm/fbackupfiles/dates file.

```
# fbackup -f /dev/rmt/0m -i . -e ./subdir
```

Backup everything under the current directory except the subdir to /dev/rmt/0m.



173

## BU-10. SLIDE: The frecover Command

121

### The frecover Command

Syntax:

```
frecover -r | -x [-g graph] [-f device]
frecover -I path | -R path [-f device]
```

| | |
|---|---|
| `-r` | Recover everything that is on a backup volume. |
| `-x` | Extract certain files from a backup volume. |
| `-I` *path* | Read the index from the backup volume and write it to *path*. This retrieves a table of contents. |
| `-R` *path* | Restart an interrupt recovery |

Examples:

```
# frecover -rf /dev/rmt/0m
# frecover -x  -g /usr/adm/fbackupfiles/graphfile -f /dev/rmt/0m
# frecover -I /tmp/index -f /dev/rmt/0m
```

-f

tcio -i | update.SAC | frecover -f-

cpio -i

51436 /51482 BU-10. 10 

## Student Notes

The real reason for employing a backup strategy is to enable the recovery of lost files. As mentioned earlier, lost files can be the result of inadvertent removal by a user, or some kind of file system disaster. The **frecover** command is the counterpart to the **fbackup** command. It is designed to retrieve files from backups that were created with **fbackup**. Like **fbackup**, **frecover** is flexible and has many options that modify its default mode of operation. Only a few key options are explained here. The **frecover(1M)** manual page contains a full description of all the options.

There are four basic modes of operation for **frecover**:

- **frecover -r**
  Recover everything that is on a backup volume.
- **frecover -x**
  Extract certain files from a backup volume.

- **frecover -I** *path*
  Read the index from the backup volume and write it to *path*. This retrieves a table of contents.
- **frecover -R** *path*
  Restart an interrupted recovery. The option to restart an interrupted recovery will not be discussed here.

Unlike **fbackup**, **frecover** does have a default for input. The default is **/dev/rmt/0m**. If a different input source (device) is to be used, a **-f** *device* option may be specified on the command line. As with the **fbackup** command, **-f -** can be used to specify that the standard input should be used.

**frecover -r** and **frecover -x** have some options in common. A few of them are **hoFX**:

```
frecover -x | -r [ -hoFX ]
```

The **-h** option is used to recover (or extract) only directories and not the files contained in them.

The **-o** option is used to force **frecover** to overwrite a newer file with an older one. Normally, **frecover** will not overwrite an existing disk file with an older version of the file.

The **-F** option causes **frecover** to strip all the leading directories from the path names of files being recovered. If **/usr/bin/vi** and **/bin/sh** were on the backup and were recovered using the **-F** option and the current working directory were **/users/root**, the resulting files would be **/users/root/vi** and **/users/root/sh**.

The **-X** option makes all recovered files relative to the current working directory. Suppose the current working directory was **/users/root** and the file **/usr/bin/vi** were being recovered. With the **-X** option, the file would be deposited in **/users/root/usr/bin/vi**. This option can be very useful when you are unsure about the directory and files that might result from an **frecover** session.

An option that is unique to the **frecover -x** mode is **-g** *graph*. This allows you to use a graph file in the same way as with **fbackup**. The file format for the *graph* file is the same. Lines that begin with an **i** indicate a path that is to be included in the recovery. Lines that begin with an **e** indicate a path that is to be excluded from the recovery. This is useful for partial recoveries.

The first example on the slide recovers all files from high density reel tape. The second example recovers all files indicated in graphfile. The third example retrieves an index of files from the tape and puts it in **/tmp/index**.



174

## BU-11. SLIDE: The dd Command

122

### The dd Command

Syntax:

dd [ *option=value* ] . . .

Options:

| | |
|---|---|
| `if=`*file* | Input file name |
| `of=`*file* | Output file name |
| `bs=`*n* | Set both input and output block size to the same size |
| `count=`*n* | Copy only n input blocks |
| `conv=`*value* | Data conversion option. *value* could equal `ascii, ebcdic, ibm, lcase, ucase`, etc. |
| `skip=`*n* | Skip *n* input blocks before starting copy |
| `seek=`*n* | Seek *n* output blocks before starting copy |

Examples:

Disk to disk copy

```
# dd if=/dev/rdsk/c1d0s2 of=/dev/rdsk/c2d0s2 bs=32k
```

Tape to disk copy of foreign tape with EBCDIC to ASCII translation

```
# dd if=/dev/rmt/0h of=/users/dave/ascii.data conv=ascii bs=80
```

51436 /51482 BU-11. 11 

## Student Notes

The dd command is useful in some limited situations, but is technically *not* a backup command.

One useful function of the dd command is to copy from a disk to a disk. In the first example, dd is being used to copy an entire disk from bus address 1 to the disk at bus address 2. The block size that is specified (512 kilobytes) is an optimum number for doing disk-to-disk copies due to the size of the buffers allocated to the disk drivers. Note that while this is very fast, it requires a destination disks the same size or larger for successful operation.

The dd command can also be used to read a foreign magnetic tape format. In the example above, dd is reading a file off the tape using a block or record size of 80 bytes and translating the contents from EBCDIC to ASCII format. The final ASCII version is stored in the file `/users/dave/ascii.data` on the disk.

Another example follows:

```
# dd if=/dev/rdsk/c1d0s9 of=/dev/rmt/0h count=4480 bs=32k
# dd if=/dev/rdsk/c1d0s9 of=/dev/rmt/0h count=4480 bs=32k skip=4480
# dd if=/dev/rdsk/c1d0s9 of=/dev/rmt/0h count=4480 bs=32k skip=8960
```

This example show backing up a large disk section to a mag tape. The disk input is too large to fit on a single reel, so a separate dd command must be used for each of the three reels used for the backup. To restore this dd backup, execute the following:

```
# dd if=/dev/rmt/0h of=/dev/rdsk/c1d0s9 count=4480 bs=32k
# dd if=/dev/rmt/0h of=/dev/rdsk/c1d0s9 count=4480 bs=32k seek=4480
# dd if=/dev/rmt/0h of=/dev/rdsk/c1d0s9 count=4480 bs=32k seek=8960
```

The dd command is *not* recommended for system backups because:

- It cannot do selective restores
- Since a direct image copy is produced, good blocks as well as bad blocks are saved

## BU-12. SLIDE: Using SAM to Backup and Recover Files

123

### Using SAM to backup and recover files

Syntax:

```
SAM                    System Administration Manager

           Highlight an item and then press "Return" or "Select Item".

                     Users ->
                     Groups ->
                     Auditing and Security (Trusted System) ->

                     File Systems ->
                     Peripheral Devices ->
                     Backup and Recovery ->

                     Networks/Communications ->
                     Cluster Configuration ->
                     Kernel Configuration ->

                     Task Customization  . . .
                     Other Utilities ->

                     How to Use SAM

1 Help  2  3 Shell  4 Select Item      5      6      7      8 Exit SAM
```

51436/51482 BU-12. 12 

## Student Notes

The System Administration Manager utility is useful in most backup situations. SAM now supports the **fbackup** and the **frecover** commands for all HP-UX systems. These commands *are* recommended for system backups.

The screen depicted in the slide, allows us to enter the SAM backup and recovery facility.

## BU-13. SLIDE: SAM: Backup and Recovery

124

### Using SAM: Backup and Recovery screen

```
SAM                          Backup and Recovery

         Highlight an item and then press "Return" or "Select Item".

            Backup Files Interactively ...
            Add an Entry to the Automated Backup Schedule ...
            Remove an Entry from the Automated Backup Schedule ...
            View Currently Scheduled Backups

            Recover Files or Directories ...
            Get List of Files from a Backup Tape ...

            Show Backup and Recovery History
            View Information on Last:
                Full Backup
                Incremental Backup
                File Recovery

1 Help  2 Main Menu  3 Shell  4 Select Item  5  6  7  8 Previous Menu
```

51436/51482 BU -13. 13 

## Student Notes

The System Administration Manager will take you to this screen if you selected "Backup and Recovery" in the main menu. The screen depicted in the slide, allows us to select a variety of backup and recovery functions.

- Backup Files Interactively
- Add an Entry to the Automated Backup Schedule
- Remove an Entry from the Automated Backup Schedule
- View Currently Scheduled Backups
- Recover Files or Directories
- Get List of Files from a Backup Tape
- Show Backup and Recovery History
- View Information on Last Full Backup

- View Information on Last Incremental Backup
- View Information on Last File Recovery

Let's look at the screen produced by selecting the menu item: "Backup Files Interactively".

## BU-14. SLIDE: SAM: Backup Files Interactively

125

### SAM: Backup Files Interactively

```
SAM                    Backup Files Interactively

      Fill in or modify the desired fields and then press "Perform Task".

         Files to be Included                 Files to be Excluded
              /users

   Device file:   /dev/update.src

   Create index file:  y

                        —working—

1 Help  2 Main Menu  3 Shell  4 Perform Task  5  6  7  8 Exit Task
```

51436/51482 BU-14. 14 

## Student Notes

In this screen you would list the directories you wished to selectively backup, and any sub-directories you want to exclude. You have the opportunity to specify the device file you want to use. The Help screen will tell you what backup devices are available on your system.

If you fill in the form and then select **Perform Task** you will first see a brief message:

```
-- working --
```

which will flash at the bottom of the screen for a few seconds. Afterward, a warning will be posted in a pop up box to confirm your intent to perform an interactive backup:

```
+--------------------------------------------------------------------------+
|                                                                          |
|NOTICE:  This backup operation may take some time.  The SAM interface will|
|  disappear and the HP-UX command "fbackup" will execute.                 |
```



137

```
|                                                                      |
| Please check that a tape is loaded and the backup device is ready before|
| proceeding.                                                          |
|                                                                      |
| Are you sure you want to proceed? (y or n)                           |
|                                                                      |
+----------------------------------------------------------------------+
```

At this point you should be sure that a suitable backup medium is placed in the backup device. It should be ready to use before you continue. You should select y to continue the backup process. After you do this, the SAM interface will disappear and you will see messages posted to the console starting at the top of the screen:

```
Starting interactive backup ...
fbackup(1421): no history is available for graph file /usr/sam/backup/grapha067
(below level 0)
fbackup(1004): session begins on Tue Feb 19 19:10:06 1991
fbackup(3212): writing volume 1 to the output file -
    1: /
    2: /users
    3: /users/rjb
    4: /users/rjb/sample
    5: /users/rjb/test
    6: /users/rjb/test/case
    7: /users/rjb/test/case/file1
    8: /users/rjb/test/case/file2
    9: /users/rjb/test/filea
   10: /users/rjb/test/fileb
        .
        .
        .
fbackup(1005): run time: 21 seconds
```

After the backup completes, another flashing message will indicate you have finished:

```
Press space bar to continue...
```

At this point you can return to SAM, and by successive presses of F8 to return to your session, or perform other backups in the SAM utility.

## BU-15. LAB: Backing Up the System (Series 300/400/700)

### Directions

Work with your partner to perform the following exercises.

Before beginning this lab, change your working directory to the /tmp directory and make three directories: cpio_dir, tar_dir, and fbackup_dir.

1. Backup the directory /usr/lib/term to a suitable backup media using the find and cpio commands. Restore /usr/lib/term from the backup to /tmp/cpio_dir.

2. Use tar to create an archive of all files under /usr/lib/term. Recover /usr/lib/term from the backup to /tmp/tar_dir.

3. Use fbackup to backup your home directory and everything underneath to a suitable backup media. Now recover one file from the backup to /tmp/fbackup_dir on your root disk.

4. Get a table of contents for all three backups and store them in a file called contents in the appropriate directory.

5. Backup the directory /usr/lib/term to a suitable backup media using the SAM utility. Create an index file in your home directory.



172

## BU-16. LAB: Backing Up the System (Series 600/800)

### Directions

For this lab you will simulate copying files to a magnetic tape drive. This will allow everyone in the class to perform backups simultaneously without worrying about contention for peripherals. Create a directory called `/tmp/`*your_login_name* with a file called `magtape` underneath. The file `/users/`*your_login_name*`/magtape` looks like a very fast magnetic tape device.

Remember that some commands care whether you are naming a file or a device file as a backup destination.

Now create the following directories: `$HOME/cpio_files`, `$HOME/tar_files`, and `$HOME/fbackup_files`. These will be used to restore your files from tape.

(Hint: Remember that the simulated magtape referred to in the lab questions is `/tmp/`*your_login_name*`/magtape`.

1. Backup all the files in your `$HOME` directory to your simulated magnetic tape device using the `fbackup` utility. Check to see that the files were actually copied by using the `frecover` command to create an index file called `/tmp/`*your_login_name*`/index`.

Now use `frecover` to retrieve your files and put them in the directory `$HOME/fbackup`. Check to see that the files were restored properly. Then remove these files.

2. Use the `find` and `cpio` commands to backup the directory `/usr/lib/term` to your simulated magtape. Now look at the "tape" table of contents. Restore these files from magtape to the directory `$HOME/cpio_files`. Verify that the restoration was successful. Then remove these files. *Remember to use the relative (not absolute) form of the find command when performing the copy.*

3. Use the `tar` command to backup the directory `/usr/lib/term` to your simulated magtape. Now look at the "tape" table of contents. Then restore these files from magtape to the directory `$HOME/tar_files`. Remember once again to use the *relative form* of the `tar` command and to remove the restored files when you are done.

4. Backup the directory `/usr/lib/term` to an actual peripheral on the system such as cartridge tape, mag tape, or DDS format DAT using the SAM utility. Create an index file in your home directory.

Check with your instructor for specifics regarding the peripheral to ensure there is no conflict with shared peripherals on your system.

# Module PI — Post Installation

## Objectives

Upon completion of this module, the student will be able to:

- Explain those tasks that should be performed immediately after completing a system installation.
- Set a superuser password.
- Remove optional filesets from the system using `rmfn(1m)`.
- Format on-line manual pages.
- For a Series 300/400/700 system, make and use a recovery system.

# Module PI — Post Installation

## PI-1. SLIDE: Post Installation Procedures

126

### Post Installation Procedures

- Execute basic procedures to ensure system integrity.
- Customize the system
- Create a recovery system on a Series 300/400/700

51436 /51482 PI-1. 1 

## Student Notes

Once installation is finished, the job of system administration is just beginning a new facet. But before we go onto the daily care of an HP-UX system there are some steps we need to take:

- Execute basic procedures to ensure system integrity.
  In order to guard against sabotage and recover from failures, there are some safeguards that can be employed. We'll look at these next.
- Customize the system
  The system needs to be customized for such things as its name and time zone. Files executed at login will need modification. Unwanted files will need to be removed, and man pages will need to be formatted.
- Create a recovery system on a Series 300/400/700
  A recovery system needs to be created on a Series 300/400/700. A recovery system is used in case the system is unbootable due to a corrupt root disk. A recovery system provides tools to repair a corrupt root disk. On a Series 600/800 the support tape is used in such situations.



120

## PI-2. SLIDE: Procedures to Ensure System Integrity

127

### Procedures to Ensure System Integrity

- Unload all *UPDATE* tapes and store in a safe place
- Assign a root password
- Check install/update information in `/tmp` and `/etc/newconfig`
- Print superblock information from `/etc/super_blocks` or `/etc/sbtabs`

51436 /51482 PI-2. 2 

## Student Notes

### Unload all *UPDATE* tapes and store in a safe place

As soon as possible following the installation, you should unload all the *UPDATE* media you used in the installation, and along with *INSTALL* media, store them in a safe place. Failure to do this may result in a security problem. A person in possession of these tapes who has physical access to your system can destroy data and compromise your system security.

### Assign a Root Password

To ensure that your system is secure, assign a root password immediately. Use the **passwd** command to set a superuser password, as shown below. The passwords you type will not be echoed to the terminal, and you must enter them twice.

```
# passwd root
```

```
Changing password for root
New password: ********
Re-enter new password: ********
#
```

You should not write this password down anywhere, as it would be more likely to be discovered by someone else. However, you should forget this password, there is no easy way to change it. On a Series 600/700/800 system, you must boot in single user mode to change the password. (See Boot Up (Series 700) module and Installing HP-UX (Series 600/800) for details for the Series 700 and Series 600/800 respectively.) On a Series 300/400 and also a Series 700, the recovery system can be booted and the passwd then changed. (See material at the end of this module for details.)

### Check install/update information

After the installation you must check the install/update information logged during the process. To do this you should check the log files in the `/tmp` directory and check files placed in the `/etc/newconfig` directory.

The log file `/tmp/update.log` recorded a log of any errors or warnings that happened during the install. A careful reading of that file is important to determine if the installation was correctly carried out. The entries entered in this file were discussed in the module on "Installing HP-UX", which you can consult for the meaning associated with errors, warnings, and "Notes" posted to the log.

The `/etc/newconfig` directory contains some files that hold information about your new release. The file `/etc/newconfig/`*UpdateInfo* gives you information on the release you are installing. A careful reading of this file is important prior to performing an update, and less critical to an installed system, unless you need to know what changed since the previous release of HP-UX.

The `/etc/newconfig` directory contains files that were referenced in messages from the update.log file mentioned above. These are usually copies of files that you will need to copy into another directory, based on update.log "Note" messages.

### Print Superblock Information

It is important that you know where the backup copies of your primary super block are located. On a Series 600/800 the file `/etc/super_blocks` contains this information. On a Series 300/400/700, the file `/etc/sbtab` contains this information. Print a copy of the appropriate file for your system, and save the printout where it can be retrieved. A good place for keeping this important data is in the binder you use for installation and maintenance history and for a system log.



181

## PI-3. SLIDE: Checking up on update

128

### Checking up on "update"

- Verify time is correct by examining the contents of `/etc/src.sh` and `/etc/src.csh`
- Remove files not removed during `update`

51436 /51482 PI-3. 3 

## Student Notes

### Verify time is correct

Several startup customization files are "sourced" to customize the system configuration at boot time (set the time zone and host name). You may choose to verify their contents at this time; however, at HP-UX 8.0 these are edited for you as part of the menu that comes up after installation utility has installed the software.

To verify customized data in startup files, check two files created at the end of the installation process, **`/etc/src.sh`** and **`/etc/src.csh`**. As of the HP-UX 8.0 release, you no longer have to edit certain system files to ensure that the time zone environment variable (TZ) and the system hostname are correctly set. TZ is referenced in

- **`/etc/rc`**
- **`/etc/profile`**
- **`/etc/csh.login`**

- `/etc/powerfail` on the Series 600/800

The two files `/etc/src.sh` and `/etc/src.csh` were created in the installation process to set the TZ variable and hostname for you at startup time. They are now "sourced" by the `/etc/rc` script at boot time and the system-wide profiles (`/etc/profile` and `/etc/csh.login`) as users login.

The contents of `/etc/src.sh` should be something like:

```
TZ=CST6CDT; export TZ
SYSTEM_NAME=training ; export SYSTEM_NAME
```

The contents of `/etc/src.csh` should be something like:

```
setenv TZ CST6CDT
```

The time zone and host name values will obviously vary depending upon your location. The values here are typical for the central USA timezone. To recreate these two files you can re-run the interactive program `/etc/set_parms`.

The files `/etc/src.sh` and `/etc/src.csh` are "sourced" inside shell scripts at startup time to set the environment variables used in `/etc/rc` for setting the network name of the system and the time zone for formatting the `date` command output.

The program `/etc/set_parms` will not reset the system clock for you if the file `/etc/installtime` exits. You should not delete the file `/etc/installtime` as its time stamp is a reference indicating when you system was installed.

For more information on time zones see the file `/usr/lib/tztab`.

### Remove files not removed during update

During the `update` portion of the installation, the install or update process may be unable to remove files from the file system. This occurs when there is an attempt to remove a file from the disk when a process is running that corresponds to the file. As an example, it may well be that the file to be removed was `/etc/update`, and it was also executing. the program would rename it to `/etc/#update`, and unload a new `/etc/update` from the update media.

The `update` process will usually rename that file with a name formed by placing an "#" in front of the old filename. You should remove all such files after the installation or update completes. The list of such files is in `/tmp/update.cleanup` and may be removed with the command:

```
# rm -rf `cat /tmp/update.cleanup`
```



192

## PI-4. SLIDE: Startup Files that Customize Your Configuration

129

### Startup Files that Customize Your Configuration

The superuser profile -
`/.profile`

System-wide user profiles -
`/etc/profile` and `/etc/csh.login` Worldwide

Default user profiles -
`/etc/d.profile`, `/etc/d.login`, and `/etc/d.cshrc` User SAM

System initialization script -
`/etc/rc`

local RC

51436 /51482 PI-4. 4 

## Student Notes

Startup files will need to be customized for your configuration. These files include:

- The superuser profile - `/.profile`
- System-wide user profiles - `/etc/profile` and `/etc/csh.login`
- Default user profiles - `/etc/d.login`, and `/etc/d.cshrc`
- System initialization script - `/etc/rc`

### The superuser profile

You may want to customize the file `/.profile` , which sets up the superuser shell environment. This file as installed, has several shell functions defined that allow the superuser session to accommodate a wide range of terminal types. For systems with a system console that remains the same, these shell functions will quickly outlive their usefulness and become an annoyance. You can correct this behavior by editing

the files /.profile and /etc/ttytype. Do not worry about changing the profile since a backup copy of this file lives in /etc/newconfig/.profile in case you need to restore the original version.

Here are some hints for editing .profile. Look for the following lines (around line number 246):

```
else
    ask="yes"        # change to "no" if sure of value in /etc/ttytype
    echo "TERM value from /etc/ttytype: \"$TERM\"" # and uncomment
fi
```

If you are sure of your console type, change yes to no and uncomment the line indicated (line 247).

Next you need to edit /etc/ttytype and put your console type in the first line of the file replacing the default type if it does not match your terminal.

```
300h console
```

Some common terminal types you might use instead of "300h" would be entered as:

98550
hpa1096a
70094
vt100

While editing /.profile you should avoid the temptation to add "." (dot) to the superuser's $PATH variable, especially to the front of the $PATH variable. This creates a potential security problem.

## System-wide user profiles

You will need to modify startup files that are always referenced for all users at startup time. You need to add some text to these files for correct terminal configuration, and customization for the types of terminals you have on your system.

The files /etc/profile and /etc/csh.login are "system-wide" profiles that initially configure all user shells at startup. The Korn (ksh) and Bourne (sh) shells use /etc/profile, while the C Shell (csh) uses /etc/csh.login. In both of these files you need to add the two commands: **tset -e -k** and **tabs** to the script.

Immediately after an installation, the terminal tabs are not always set correctly which causes the backspace key to have unpredictable behavior. Also, if you try to use the **vi** editor, it may display far more lines than are actually in the files (it interprets the embedded tabs as new line characters). You can correct these problems by editing the /etc/profile and /etc/csh.login and adding the lines:

```
tset -e -k
tabs
```

The **tset** command sets the terminal characteristics and the **tabs** command resets the terminal tabs. If you then log out and then log back in, the system should be operating correctly. In this case, when **tset -e -k** is used, the erase and kill characters to their default values of KILL=Control-X and ERASE=Control-H (backspace). For more information, see *tset(1)* and the contents of the databases kept for terminals in /usr/lib/terminfo. These topics will be covered later in more detail.



183

Commands in these files should be kept to a minimum, but you may add shell variables that need to be set for everybody. An example would be an environment variable required by a software subsystem in order for it to work. You may be directed to do this by a *README* document supplied with an optional subsystem.

### Default user profiles

You may need to edit several files that are used when a new user is added to the system. These files create the default profiles and other startup configuration files that are put in the new users home directory to customize that users environment. You must be aware that changes to these "d-dot" files will affect every user added to your system after the change. All "d-dot" files in "/etc" are renamed with the "d" dropped to become hidden "dot" files in each users home directory.

The following files are defaults for the new user profiles that configure user shells at startup.

- `/etc/d.profile` - configures "ksh" and "sh" shells
- `/etc/d.login` - configures the "csh" shell
- `/etc/d.cshrc` - configures the "csh" shell

There are other "d-dot" configuration files which you may want to customize on a system wide basis such as `/etc/d.exrc` which configures the "vi" and "ex" editors. One customization that is useful for users new to "vi" is to change the line in `/etc/d.exrc` which reads:

```
set autoindent autowrite showmatch wrapmargin=0 report=1
```

to read:

```
set noautoindent showmode showmatch wrapmargin=0 report=1
```

These changes have the effect of:

- displaying an editor MODE indicator for the editor insert/replace mode, which is very helpful for users new to "vi"
- turning off automatic indenting when inserting text
- disabling the automatic writing of a file when a new file is edited, a dangerous feature for both experienced and novice users.

You might want to add your own configuration files such as `/etc/d.kshrc` for startup customization of the Korn shell.

### System initialization script

The `/etc/rc` system startup script prompts the system console for the correct time every time the system reboots. The TIMEOUT variable determines how long to wait for user interaction to set the date and time of day. If this variable is set to 0 the script does not ask each time. You might want to make this edit to the function "initialize" in `/etc/rc`.

```
TIMEOUT=0
```

The /etc/rc script has provisions for starting up processes or setting system wide configurations. If these are required there is a function called "localrc" that could be edited to add the commands. We cover the rest of the /etc/rc script later in this course.

The /etc/rc script lacks an automated way to mount CD-ROM drives that are locally connected. You might want to add a function to the script and call it after the HFS file system is mounted. A suggested function would be:

```
cdfsmount()
{       # Mount CDFS volumes listed in /etc/checklist:
    # HFS volumes are mounted via rc using hfsmount
    # NFS volumes are mounted via netlinkrc
    if /etc/mount -a -t cdfs
    then
        echo local CD-ROM file systems mounted
    fi
}
```



134

## PI-5. SLIDE: `rmfn`: Removing Optional software

130

### `rmfn`: Removing Optional software

rmfn Delete Partitions

Mark "y" or "n". To pick and choose individual filesets within a partition, press "Select Filesets". A "p" means that some filesets have been selected within a partition. Press "Start Removing" when selection is complete.

| Delete | Partition | Partition Description | Size in Kbytes |
|---|---|---|---|
| n | DIAGNOSTICS | Hardware Diagnostic Programs | 2 |
| n | GRAPHICS | Graphics Products | 12071 |
| n | NETWORKING | Networking Products | 8280 |
| n | NLS | Native Language Support | 592 |
| n | OS-ADMIN | Recommended Administration Cmds | 7980 |
| n | OS-CORE | Recommended System Core | 11022 |
| n | OS-FEATURES | Selectable OS Features | 18626 |
| n | PROG-LANGUAGES | Programming Languages | 10564 |
| n | REFERENCE-DOC | Reference Manual Pages | 5753 |
| n | SHARED-LIBS | Run-time Shared Libraries | 4322 |
| n | WINDOWS | Windowing Products | 20793 |

Help | | Shell | Start Removing | View Selected | Select Filesets | | Exit rmfn

51436/51482 PI-5. 5 

### Student Notes

It is recommended that when you do your installation, you install all products on your distribution media. After everything is installed, you can remove the filesets that you do not want. The filesets that are installed on your system are listed in the directory **`/etc/filesets`**.

Starting at the HP-UX 8.0 release, the **`rmfn`** command can be used to remove unwanted optional filesets. The **`sysrm`** command was used for this purpose previously, and is now obsolete. The **`sysrm`** command had a serious limitation: it allowed you to remove a fileset needed by another product's files. The interactive utility **`rmfn`** has a much better understanding of fileset dependencies, and will prevent this from happening accidentally.

To use the interactive utility **`rmfn(1m)`** command to remove the unwanted optional filesets invoke it with:

```
# rmfn
```

You will have to wait a short time while the filesets that are loaded on your system disks are scanned for size and dependency information. During the delay you will see the following message on the screen:

```
Getting fileset information  ...
```

Then the screen shown in the slide will appear, listing the partitions loaded on your system, along with their sizes. At this point you may select a partition for deletion by changing the **n** to a **y** in the "Delete" column. When you select a partition for removal it will again take some time for the system to traverse the list of dependencies for the filesets in the partition you select for removal. This is usually ten or twenty seconds of intense disk activity before the **y** is posted to the screen.

**Note** Do not press **y** twice, thinking that it did not recognize the first one, because this may cause you to select another partition you did not intend. The program moves the cursor to the next line after a selection is made. A second affirmative response will select the next partition.

After you have selected a partition, the softkey `View Selected` can be used to determine exactly what filesets will be removed.



185

## PI-6. SLIDE:rmfn: View Partitions/Filesets Selected

131

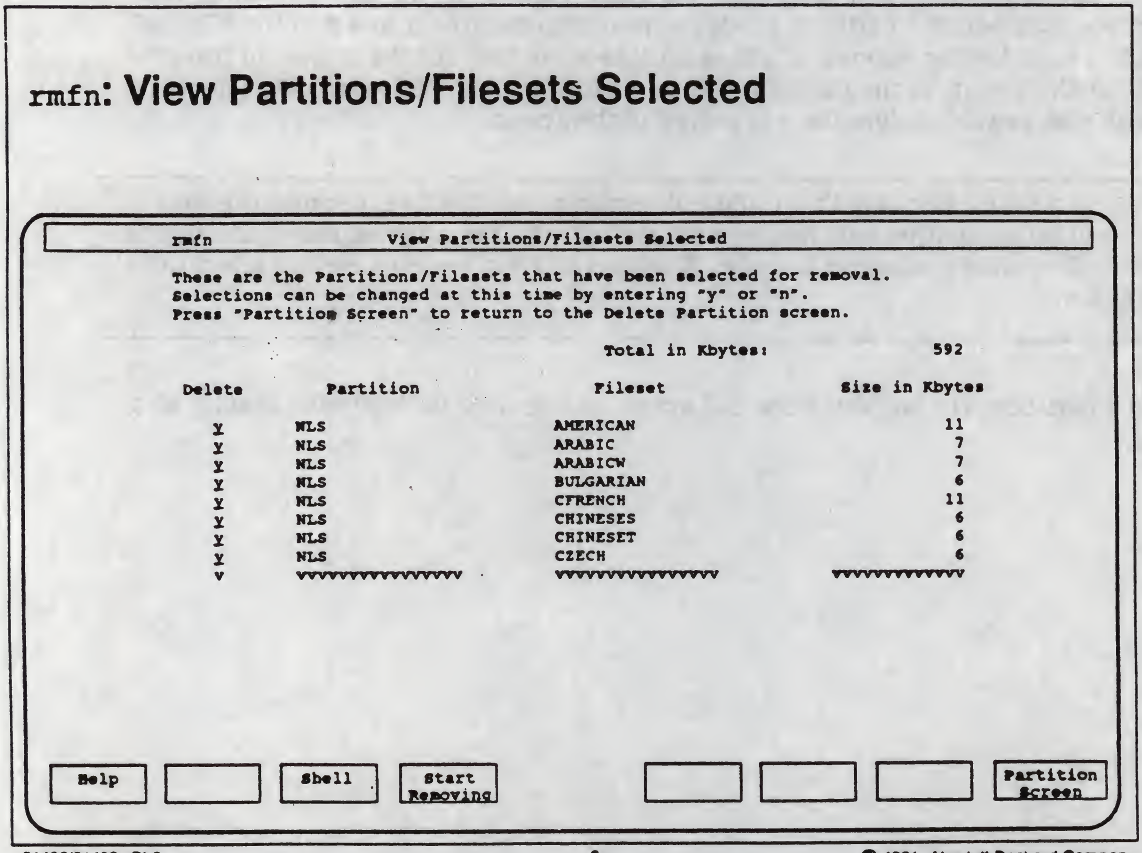


51436/51482 PI-6. 6 

### Student Notes

This is the screen you will see when you select a partition and press the softkey **View Selected**. While in this screen you may decline to delete any of the selected filesets by changing the "Delete" field (shown on the slide as **y**) to **n**. The same caution stated above applies to multiple key presses in this screen.

If a dependency exists for a fileset selected for deletion, you will get a pop-up window that warns you about it and offers you a chance to also delete the dependent fileset. You must answer **y** or **n** depending upon your requirements. If you answer negatively, neither the dependent fileset or the corresponding fileset you had marked will not be removed.

Once you have selected all the filesets you want to delete, you must press **Start Removing** softkey. You will be asked to confirm your choice to start the removal process. A pop-up window will appear. No softkeys will be defined, but the empty softkey labels will be displayed. The only possible response is **y** or **n** at this point. To confirm that you want to go on, press"**y**".

Of course the number and sizes of files you choose will affect the number of kilobytes to be removed. This is you last chance to back out. You must answer y to start removing files. Once you have confirmed your choice to start the removal process, a new window will appear to provide a status report on the removal process.



186

## PI-7. SLIDE: rmfn: Removing Filesets

132

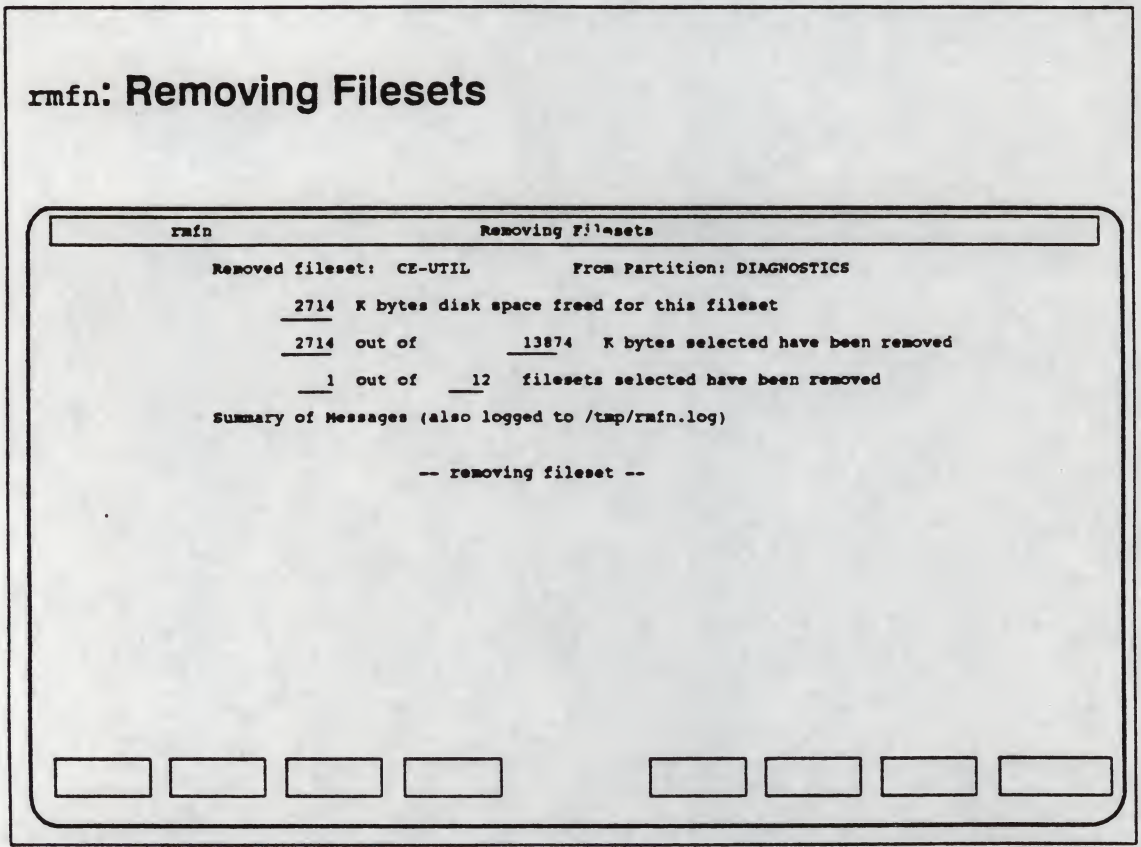


51436/51482 PI-7. 7 

## Student Notes

This screen presents you with a continuous status report about the process of fileset and partition removal. It will display the progress and any error messages until the process completes. No softkeys are displayed while this screen appears.

The removal may take several minutes as individual filesets are removed. You will see a message flash at the bottom of the screen:

```
-- removing fileset --
```

At the end of the removal you will see a screen to indicate the end of the removal process. This screen will also report the disk space freed for this fileset. After this menu, you should select **Exit rmfn** if there are no other partitions and filesets to remove.

## PI-8. SLIDE: Format Manual Pages

132

### Format Manual Pages

You have different options for preparing on-line documentation:

- Create formatted versions of all manual pages

  # /etc/catman

- Create formatted versions of some sections of manual pages

  # /etc/catman 1 1m 2

- Do not create any formatted manual pages, but instead "build-as-you-use"

  # man somecmd

51436 /51482 PI-8. 8 

## Student Notes

Before we explain **catman** it helps if you understand how the manual entries are stored on the system. All the entries, both unformatted (source) and formatted, are stored in sub-directories under the **/usr/man** directory. Both unformatted and formatted entries can appear in compressed or uncompressed form. This gives us four possibilities.

**Table PI-1.**

| | unformatted | formatted |
|---|---|---|
| **compressed** | /usr/man/man*X*.Z/* | /usr/man/cat*X*.Z/* |
| **uncompressed** | /usr/man/man*X*/* | /usr/man/cat*X*/* |

*X* is the section: 1, 1m, 2, 3, 4, 5, 7, 9.
The unformatted entry is an **nroff** source.



127

**catman** (with no options) creates the formatted versions of the on-line manual pages from the compressed **nroff** source files. (Compressed files end in .Z.) When you invoke **catman**, each manual entry in the man*X*.Z directory is examined and those whose formatted versions are missing or out-of-date are created, compressed, and put it into the cat*X*.Z directory. Before invoking **catman**, it is best if you remove all of the **catX** directories (leaving only the **catX.Z** directories).

If you want to create uncompressed entries, make sure the **catX** directories exist and invoke **catman -z**. This command puts an uncompressed version of each entry in its appropriate cat*X* directory.

You can choose to format only certain sections of the manual pages by invoking **catman** *section*, where *section* is the number of the section you wish to format, for example sections 1, 1m and 2 as shown by:

```
# /etc/catman 1 1m 2
```

There are several user commands that have names used to reference the manual pages that appear in more than one section of the manual. Separate pages may exist for a user command (section 1), system administration command (section 1m), library routine (section 3), system call (section 2), file format (section 4), special file (section 7), or miscellaneous entry (section 5). Several examples are: **mkdir(1)**, **mkdir(2)**, **mknod(1m)**, **mknod(2)**, **mknod(4)**, **mount(1m)**, **mount(2)**, and **mount(3)**.

To read a specific manual page a user must type the **man** command with the section number as the first parameter. Several examples follow:

```
$ man 1 mkdir  OR $ man 2 mkdir
```

You can choose not to format any manual entries with **catman**. If you do not format entries, then the first time you try and access a manual entry you will get this message:

```
Reformatting entry.  Wait...
```

You will wait a short period of time before the manual entry appears on your screen. This is usually less than 10 seconds, but may be much longer for very large manual pages such as "sh". After the first time the page is read with **man some-command**, the entry for **some-command** is added to the appropriate cat*X* directory, and in subsequent accesses, the user does not have to wait. This is a "build-as-you-use" alternative. The system fills the cat*X* directories as the users access commands with **man**. For this to work, the cat*X* directories must exist. If they do not, you can create them with the following script:

```
cd /usr/man
for num in 1 1m 2 3 4 5 7 9
do
  mkdir cat$num
done
```

**Note**

If you execute **catman** it requires several megabytes of disk space and several hours to complete. By doing so you get improved performance, but you sacrifice disk space. (To recover some disk space, you could delete the unformatted directories after using **catman**.)

Running **catman** allows you to use the **man -k keyword** command to locate documentation, executables and libraries by keyword lookup. A reference to the keyword in the one line description of the command is used as a search key for the "/usr/lib/whatis" database built by the **catman** command. All commands (files, man pages) with that keyword are also listed. This is like a "conceptual" cross reference to commands based on the keyword used.



128

## PI-9. SLIDE: Making a Series 300/400/700 Recovery System

133

### Making a Series 300/400/700 Recovery System

A recovery system is useful if:

- The root password is unknown (or /etc/passwd was accidentally deleted)
- You have a bad boot area on your system disk
- There is no login prompt running on the console

To make recovery system on:

| | |
|---|---|
| Cartridge tape – | # mkrs -v -t ct |
| SCSI DDS/DAT tape – | # mkrs -v -t dat |
| Optical disk – | # mkrs -v -t od |

51436/51482 PI-9. 9 

## Student Notes

You should make a recovery system every time you install or update your system. Making a recovery system does not take much time, and it could save your system's life. Once you have made it, however, ***lock it up*** because it could be used to break into your system. (Note, you must have the **SYS-ADMIN** fileset to create a recovery system at 8.0 or later. The command **/etc/mkrs** is in this fileset.)

As you can see from the slide you can make your recovery system on a variety of devices. If you choose to make your recovery system on cartridge tape, use a 150-foot cartridge tape. You could use a 600-foot tape, but it takes much longer. It takes 1-2 hours to make a recovery system on a 150-foot tape, while it takes approximately 6 hours to make it on a 600-foot tape.

To make a recovery system, use the **mkrs** command:

**mkrs [-v] [-q] [-s] [-f** *rcdev*] **[-r** *rootdev*] **[-m** *series*]

where:

- `-v` is the verbose option.
  This is recommended, to make you feel comfortable about the progress of the creation of the recovery system.
- `-q` is the quick option.
  If enough free disk space is available in `/usr/tmp` (typically 10-20Mb), the `-q` option can be used to make *mkrs* create an image of the recovery system in this directory before copying it to the recovery media. This option generally saves a great deal of time due to reduced seeking on non-random-access recovery media (cartridge tape and DAT). Note: for DAT tape recovery systems, the `-q` option is assumed.
- `-s` is the "small" option.
  This is used to minimize the size of the recovery system. This option causes some files to not be copied to the recovery system. When creating a DAT recovery system for a small memory system (8Mb or less), this option should be used.
- `-f` *rcdev* is the device on which you want to make the recovery system.
  Replace *rcdev* with the character device file. `mkrs` searches for each of the following character device files to use as the default:
  1. `/dev/update.src`
  2. `/dev/rct/c0`
  3. `/dev/rct`
- `-r` *rootdev* is the root device for which you want to make the recovery system.
  Replace *rootdev* with the block device file. `mkrs` searches for each of the following block device files to use as the default:
  1. `/dev/dsk/0s0`
  2. `/dev/root`
  3. `/dev/hd`
- `-m` *series* is the type of machine running this software, for example `-m` 300, `-m` 400, `-m`700.
  Normally, `mkrs` properly identifies the machine type. Use this option if `mkrs` is unable to identify the machine type. The Series 600/800 families of HP 9000 computers do not support `mkrs` recovery systems.

Notice that only the DAT/DDS tape device with a SCSI interface is supported by `mkrs`. The HP-IB version of DAT/DDS is not supported for creating or using a recovery tape at the 8.0 release.

---

**Note** If you make a recovery system and you later change your swap space or kernel, you must create a new recovery system.

---

The files loaded to your recovery system will vary slightly depending upon the command line options used. The file `/etc/mkrs.data` describes which files appear on which media. The following tables list the files that are present on a maximum sized recovery system (such as a DDS/DAT medium):



189

**Table PI-2.**

| /bin | /usr/bin | /etc | /lib |
|---|---|---|---|
| cat, cp, cpio, | lifcp, | boot, fsck, group, | dld.sl, |
| date, dd, echo, | tcio, | init, inittab, mknod, | libcurses.sl |
| find, ll, ln, | mediainit | mnttab, mount, passwd, | libc.sl |
| ls, mkdir, mv, | | profile, rc, reboot, | |
| pwd, rm, sh, | | sbtab, rmfn, umount, | |
| -sh, stty, | | unlink, update, utmp | |
| sync, tar | | disktab, mkrs, newfs | |
| | | recovery.tool, | |
| | | -recovery.tool | |
| | | mkrs.init, mkfs | |
| | | pre&_init&_rc, | |
| | | clusterconf | |

**Table PI-3.**

| /dev | / | /disk | /tmp | /lost+found |
|---|---|---|---|---|
| console, | hp-ux*, | (directory | (directory) | (directory) |
| dsk/root, | hp-ux.min* | used as | | |
| rdsk/root, | | mount | | |
| dsk/realroot, | | point) | | |
| rdsk/realroot, | | | | |
| dsk/real.root, | | | | |
| rdsk/real.root, | | | | |
| syscon, | | | | |
| systty | | | | |

Certain files not loaded to your recovery system if you use the **-s** command line option. The space savings allows you to use a recovery medium with lower capacity, such as a small hard disk. The files not loaded with **-s** are:

- /bin/date
- /bin/echo
- /bin/find
- /bin/ln
- /bin/pwd

- /bin/tar
- /etc/rmfn
- /etc/disktab
- /etc/mkfs
- /etc/newfs
- /etc/update
- /usr/bin/lifcp
- /usr/bin/tcio



190

## PI-10. SLIDE:Creating a Series 300/400/700 Recovery System

134

### Creating a Series 300/400/700 Recovery System

Weg'

480

```
# mkrs -v -s -q -t ct -m 300
Building a series 300 tape recovery system on /dev/update.src.
Are you sure you want to continue (y/n)? n mkrs aborted.

# mkrs -v -s -q -t ct -m 300
Building a Series 300 tape recovery system on /dev/update.src
Are you sure you want to continue (y/n)? y
Place media in drive, once busy light remains off hit <return&>. Return

 -- press Return

Creating /etc/passwd for recovery system
Creating /etc/inittab for recovery system
Creating /etc/profile for recovery system
Creating /etc/rc for recovery system
Creating device files for recovery system
  .
  .
  .
```

51436 /51482 PI-10. 10 

## Student Notes

In this section we want to show what you will see when the **mkrs** command runs on your system console.

When creating a recovery system you will have to interact with **mkrs** only at the beginning of the process. On the screen we see first a decision not to continue. Later we run the command to ask for a "quick", "small" recovery system to be created on cartridge tape on a series 300 computer system. The "verbose" option is used to get console messages about the progress of the recovery system creation. The following log of screen messages traces the execution of **mkrs** for a short recovery system. This took about fifteen minutes to complete.

```
# mkrs -v -s -q -t ct -m 300
Building a series 300 tape recovery system on
/dev/update.src.
Are you sure you want to continue (y/n)? y
Place media in drive, once busy light remains off hit
<return&>. Return
```

```
Creating /etc/passwd for recovery system
Creating /etc/inittab for recovery system
Creating /etc/profile for recovery system
Creating /etc/rc for recovery system
Creating device files for recovery system
Creating boot kernel for recovery system
        rm -f conf.o
Compiling conf.c ...
        if [ -x /bin/cc  ]; then                \
            cc="/bin/cc ";                      \
        else                            \
            cc=/etc/conf/kerncc;        \
        fi;                             \
        $cc +M -Wc,-Nd3500,-Ns7000 -Wp,-H150000 -I. - Dhp9000s200
-D&_&_hp9000s200 -D&_&_hp9000s300 -Dhpux -D&_HPUX&_SOURCE - D&_KERNEL -DKERNEL
-Uvax -DHFS -DMC68030 -DPSTAT -DSAVECORE&_300   -DREGION - DKVM -DGENESIS
-DAUTOCHANGER -DEISA -DWRITE&_GUARD&_7&_0 -DFSD&_KI -DRTPRIO -DPROCESSLOCK
-DEISA -c conf.c
        rm -f hp-ux
        ar x /etc/conf/libkreq.a exceptions.o locore.o
vers.o name.o
funcentry.o cdfs&_hooks.o
Loading hp-ux...
        if [ -x /bin/ld ]; then         \
            ld="/bin/ld";                       \
        else                            \
            ld=/etc/conf/kernld;        \
        fi;                             \
        $ld -n -o hp-ux -e &_start -x    \
                exceptions.o locore.o vers.o conf.o name.o
funcentry.o
cdfs&_hooks.o /etc/conf/libkreq.a  /etc/conf/libdreq.a
/etc/conf/libpm.a
/etc/conf/libvm.a  /etc/conf/libsysV.a  /etc/conf/libmin.a
/etc/conf/libdevelop.a  /etc/conf/libdil&_srm.a
/etc/conf/libkern.a
/etc/conf/libk.a
        rm -f exceptions.o locore.o vers.o name.o
funcentry.o cdfs&_hooks.o
        chmod 755 hp-ux
Creating proto file for recovery system
Creating recovery system image in /usr/tmp/work216/fs&_image
Warning: 371 sector(s) in last cylinder unallocated
/usr/tmp/work216/fs&_image: 2701 sectors in 6 cylinders of
16 tracks, 32 sectors
        2.8Mb in 1 cyl groups (16 c/g, 8.39Mb/g, 1216 i/g)
super-block backups (for fsck -b#) at:
 16,
42+1 records in
42+1 records out
```



19(

```
Recovery system complete
#
```

Notice that the lines above may have been folded to fit on the printed page and are not an exact reproduction of the screens. They were taken from the screen of an actual session with `mkrs` creating a recovery system on a cartridge tape.

## PI-11. SLIDE:Using a Series 300/400/700 Recovery System: Unknown Root Password

135

### Using a Series 300/400/700 Recovery System:

### Unknown Root Password

1. Power off your system (there is no alternative).
2. Put your recovery system on line and make ready.
3. Boot in attended mode and select the recovery system device.

Wait for root file system to be checked for inconsistencies.

When the following menu appears, choose option 1.

```
Select one of the following options by number:

1) Remove the root password
2) Work in a shell to perform recovery manually
3) Perform an automatic recovery
4) Exit recovery system and reboot root file system
5) Help
Selection >> 1
```

51436/51482 PI-11. 11 

## Student Notes

The recovery system will take a few minutes to boot if your recovery system media type is cartridge tape. You will have to be patient, a tape device as a file system is SLOW. The recovery system will boot from the selected media unit, then perform a consistency check on your system disk. After this it will mount the system disk to a directory on the file system of the recovery device. When this completes it will automatically present a menu as indicated on the slide.

To recover from an unknown root password, choose option 1 from the main menu. This will remove the root password from the `/etc/passwd` file in your root file system. After the root password is removed, choose option 4 from the menu. This will exit the recovery system and reboot from the root disk. You should be able to log in as root. Remember to set a new password right away.

## PI-12. SLIDE:Using a Series 300/400/700 Recovery System: Corrupt Boot Area

136

### Using a Series 300/400/700 Recovery System:

### Corrupt Boot Area

1. Power off your system (there is no alternative).

2. Put your recovery system on line and make ready.

3. Boot in attended mode and select the recovery system device.

Wait for root file system to be checked for inconsistencies.

When the following menu appears, choose option 3.

```
Select one of the following options by number:

1) Remove the root password
2) Work in a shell to perform recovery manually
3) Perform an automatic recovery
4) Exit recovery system and reboot root file system
5) Help
Selection >> 3
```

51436/51482 PI-12. 12 

## Student Notes

Corruption in the boot area can yield an un-bootable system. To fix this problem, boot from your recovery system and choose option 3 as shown on the slide. After choosing 3, you will get this information on your screen:

```
AUTOMATIC RECOVERY

The automatic recovery option will create new versions of all files
necessary to let you reboot in single user state in the root file
system.  A list of these files can be seen in the help option.

All old versions of files will be saved in the directory /tmp/recovery.xxxx
(where xxxx is the month and day) on the root file system.  This
is available so that an old file can easily be recovered.
```

```
All actions taken during the automatic recovery will be printed on
the screen but there is no user interaction once it is started.
All actions will also be logged in /tmp/recovery.log and can be
viewed after the system reboots.

Once automatic recovery is finished, you will be prompted to continue,
and then put in the recovery system main menu.


Do you want to continue the automatic recovery? (yes or no)
```

If you answer **yes**, the automatic recovery will begin. You will see messages similar to the following:

```
AUTOMATIC RECOVERY

Copying /etc/boot to boot area on root file system.
Copied old  /disc/hp-ux to  /disc/tmp/recovery.0212/hp-ux
Creating new /hp-ux on the root file system
Copied old  /disc/lib/dld.sl to /disc/tmp/recovery.0212/dld.sl
Creating new /lib/dld.sl on the root file system
Copied old  /disc/lib/libc.sl to /disc/tmp/recovery.0212/libc.sl
Creating new /lib/libc.sl on the root file system
Copied old  /disc/etc/fsck to  /disc/tmp/recovery.0212/fsck
Creating new /etc/fsck on the root file system
Copied old  /disc/bin/cpio to  /disc/tmp/recovery.0212/cpio
Creating new /bin/cpio on the root file system
Copied old  /disc/etc/init to  /disc/tmp/recovery.0212/init
Creating new /etc/init on the root file system
Copied old  /disc/bin/sh to  /disc/tmp/recovery.0212/sh
Creating new /bin/sh on the root file system
Copied old  /disc/etc/passwd to /disc/tmp/recovery.0212/passwd
Creating new /etc/passwd on the root file system
     .
     .
     .
Existing /disc/dev/console removed.
Successfully created /disc/dev/console on root file system.
Existing /disc/dev/syscon removed.
Successfully created /disc/dev/syscon on root file system.
Existing /disc/dev/systty removed.
Successfully created /disc/dev/systty on root file system.
Existing /disc/dev/tty removed.
Successfully created /disc/dev/tty on root file system.
Existing /disc/dev/null removed.
Successfully created /disc/dev/null on root file system.
Existing /disc/dev/mem removed.
Successfully created /disc/dev/mem on root file system.
Existing /disc/dev/kmem removed.
Successfully created /disc/dev/kmem on root file system.
Successfully created /disc/dev/dswap on root file system.
```



193

```
Existing /disc/dev/dsk/0s0 removed.
Successfully created /disc/dev/dsk/0s0 on root file system.
    .
    .
    .
```

When the boot area has been fixed, you will the message:

```
AUTOMATIC RECOVERY HAS FINISHED.

The root file system may now be rebooted successfully.
By selecting the option to exit the recovery system and
reboot
the root file system, you will be rebooted in the root file
system in single user state with a working file system.
Minimal fixes have been made. They are logged in
/tmp/recovery.log

PRESS RETURN TO CONTINUE.
```

You will be returned to the recovery system menu. Choose option 4—**Exit recovery system and reboot root file system**. Your boot area should be fixed.

## PI-13. LAB: Post Installation (Series 300/400/700)

1. Unless directed differently by your instructor, set a superuser password on your system.

2. Edit the super user's profile and one other file so that your system console will be recognized automatically at login time.

3. Check the `/etc/newconfig` and `/tmp` directories for installation information. Name the specific files are you looking for?

4. Check the values automatically set for the TZ variable, date and time on your system. Which file(s) is now responsible for this?

5. Look at the superblock information for your root file system.

6. Use `rmfn` to remove a partition or fileset specified by your instructor.



199

7. Optional: Make a recovery system on the device specified by the instructor.

8. Use the recovery system to "break into" your system as if you had forgotten the root password.

## PI-14. LAB: Post Installation (Series 600/800)

1. Unless directed differently by your instructor, set a superuser password on your system.

2. Edit the superuser's profile and one other file so that your system console will be recognized automatically at login time.

3. Check the `/etc/newconfig` and `/tmp` directories for installation information. Name the specific files are you looking for?

4. Check the values automatically set for the TZ variable, date and time on your system. Which file(s) is now responsible for this?

5. Look at the superblock information for your root file system.

6. Use `rmfn` to remove a partition or fileset specified by your instructor.



195

# Module MS — Memory and Swap Management

## Objectives

Upon completion of this module, you will be able to:

- Explain the concept of physical and virtual memories as applied to a demand page virtual memory system.
- Explain the difference between physical memory, available memory, and lockable memory.
- Identify the amount of memory your system has and how much of that memory is lockable.
- Understand what swap space is and why it is necessary.
- Define various terms that describe swap types.
- Evaluate your own swap space needs.
- Set up device swap space on your system.
- Set up file system swap space on your system.
- Describe the two methods for managing swap space on your system.

## MS-1. SLIDE: Memory Management

137

### Memory Management

Demand page virtual memory manages the following:

- Physical memory
- The logical address space of each process
- Primary swap space
- Dynamic swap space

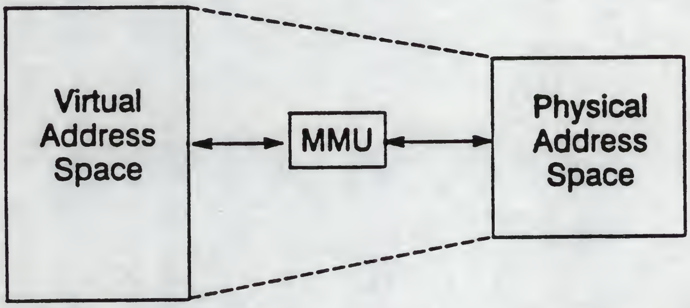


51436/51482 MS-1. 1 

## Student Notes

The term **memory management** refers to the rules that govern physical memory and allow for efficient sharing of the system's resources by user and system processes. The hardware and software that enforces these rules is collectively called the **Memory Management Unit** (MMU).

The MMU allows the total size of user processes to exceed physical memory by using an approach termed **demand-paged virtual memory**. By using the rules associated with the concept of virtual memory, when a process executes, parts of the process are brought into main memory only as needed, that is, on demand. The system uses a combination of swapping and paging to manage virtual memory. **Swapping** involves moving entire processes between RAM and mass storage (usually disk), whereas **paging** involves moving smaller units (called pages) between RAM and mass storage when needed or **on demand**.



196

For example, by using paging, a process may be written that is 4GB of code and data. This process may be executed on a system that has only 32MB of RAM. In the above example, the process' virtual or logical address space is 4GB. **Virtual address space** or logical address space is all valid memory addresses; whereas **physical address space** is the actual RAM purchased, 32 MB in the example.

In HP-UX, the basic virtual memory unit is a **page**. That is, processes are stored in blocks of memory called pages of physical memory. On a Series 600/800 computer, a page is 2048 bytes in size. On a Series 300/400/700, a page is 4096 bytes in size.

The MMU also ensures that processes do not illegally access each others' address space. Since the MMU operates at the granularity of a page, memory locking, protection, and sharing are supported at the page level.

## MS-2. SLIDE: User Process Virtual Address Space

138

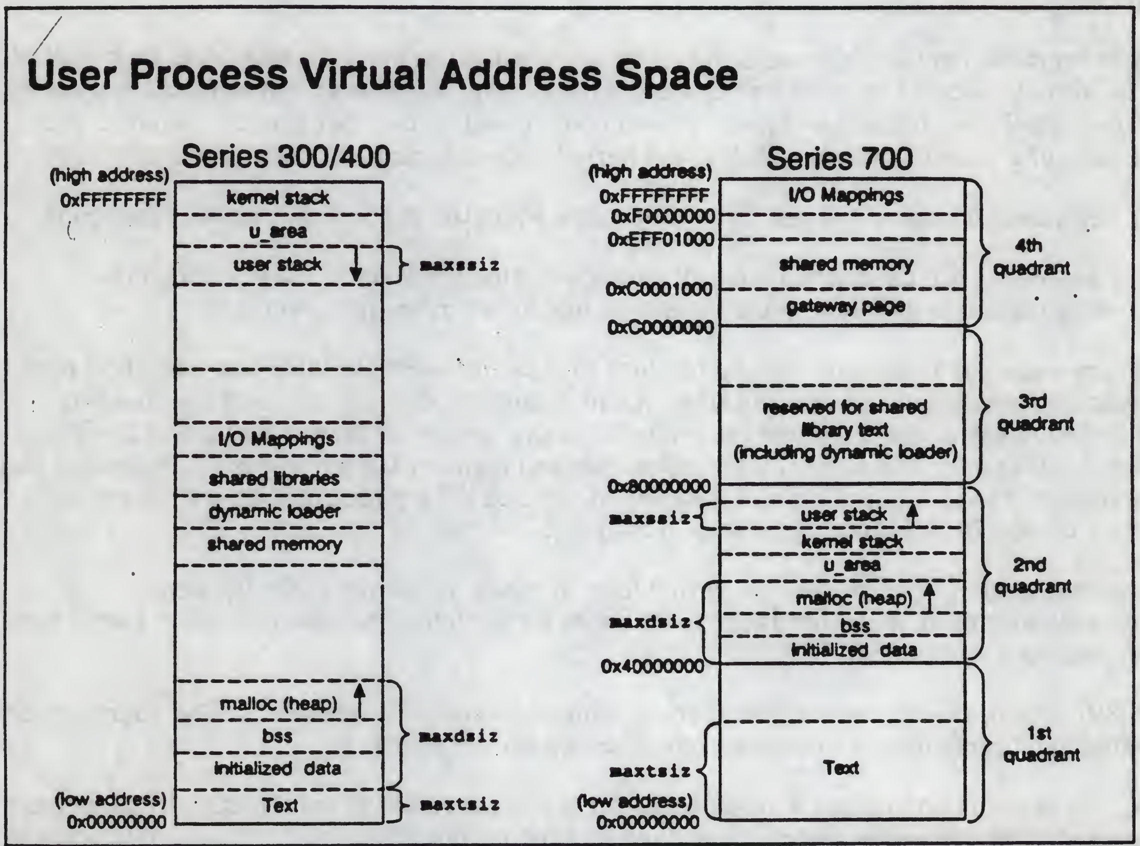


51436/51482 MS-2. 2 © 1990 Hewlett-Packard Company

## Student Notes

The HP-UX kernel imposes a virtual address space of size $2^{32}$ or 4GB on processes and indeed on the kernel itself.

The virtual address space of each HP-UX process is comprised of several regions: text (code) segment, data segment, `u_area`, kernel stack segment, user stack segment, one or more shared memory segments, and shared-library text and data segments. Each segment is divided into pages used by the demand-paged virtual memory system.

The **text segment**, also referred to as the code segment, holds a process's executable object code. It starts at low virtual address, and is followed by the data segment.



197

The **data segment** contains a process's data structures. The data segment holds initialized data, bss (uninitialized data), and heap. The heap area can be dynamically expanded into higher addresses by executing calls to `sbrk()`, `brk()`, and `malloc()` in a program's run time logic.

The `u_area` is a data structure required by the kernel when a process is executing. The `u_area` keeps track of process resources such as open files, system call returns, effective user and group ids, signal handlers, and others.

The **kernel stack segment** contains a process' run-time stack when the kernel is executing on behalf of this process. The kernel executes on behalf of processes frequently. When ever a system call is made by a process, it is the kernel that actually executes its own code to satisfy the user process' request. For example, in the case of a `time()` system call, it is the kernel code that actually reads the system clock.

The **user stack segment** contains a process's run-time stack when the process is executing user code.

**Shared memory segments** are used when multiple processes must share data (for example, in a windowing system, all window processes must be able to update a screen data structure).

The **shared library segment** is used for library routines that are dynamically loaded at execution time. Shared libraries are different from archive libraries. Archive libraries append copies of routines into a process' text and data segments. A shared library reduces the amount of memory occupied by code during execution because only one copy of the routine exists in memory for all processes. A process that uses shared libraries will have its own copy of data for the shared library routine, but will access the single shared copy of text for all processes that is in memory.

**I/O mappings** are the ranges of addresses the kernel uses to access I/O devices. On the Series 600/700/800, the **gateway page** is the hardware mechanism for switching between user and kernel mode privileges when making a system call.

On the Series 600/700/800, the virtual address space is addressed via 32-bit addresses. The address space is divided into four quadrants of one Gigabyte each as shown on the slide.

On a Series 300, the segments of virtual address space for a process reside in contiguous virtual memory. To prevent segments from overwriting each other, the system does not allow them to overlap. If system calls or stack growth do cause segments to overlap, an error occurs (such as ENOMEM) and/or the process dies.

## MS-3. SLIDE: System Memory

139

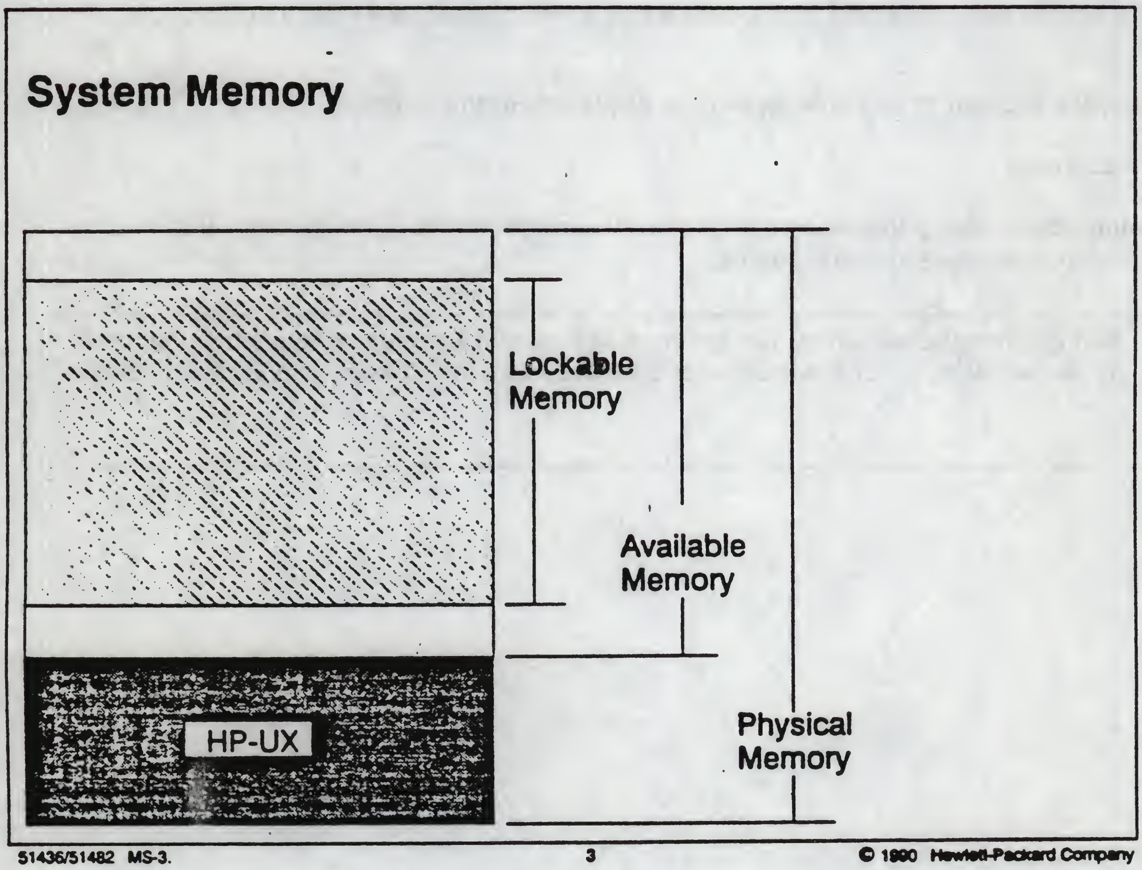


## Student Notes

**Physical memory** is the RAM (random access memory) installed in your computer. At system startup, the system displays on the system console the amount of physical memory installed:

```
real mem = xxxxxxx
```

System performance can degrade if the system spends too much time swapping - a condition known as thrashing. Remedy thrashing by adding more physical memory to your system.

Not all physical memory is available to HP-UX processes. Some memory is reserved for kernel code and data structures. The amount of memory remaining is referred to as **available memory**, and is used by the system for demand paging. During system startup, the system displays on the system console the amount of available memory:



150

```
avail mem = xxxxxxx
```

All or part of available memory can be locked by a subsystem or by user processes. **Lockable** memory cannot be swapped out to disk. Typically, locked memory holds frequently accessed programs or data structures. By keeping them memory-resident, system performance improves. If most of the available memory is locked the system may deadlock. Some unlockable memory must be available to prevent deadlock.

During system startup, the amount of lockable memory is displayed on the console:

```
lockable mem = xxxxxxx
```

Available memory minus the memory locked by subsystems or user processes is the memory that is actually available for virtual memory demand paging.

---

**Note** The `dmesg` command will show you the messages output by your system when you boot up. These messages include the amounts for **`real mem`**, **`avail mem`**, and **`lockable mem`**.

---

## MS-4. SLIDE: Swap Management

140

### Swap Management

- Size of swap is determined by system load and physical memory
- Swapping demands are monitored by `vhand`, `statdaemon`, and `swapper`
- Swap space is reserved for a process' data, bss, and shared memory segments
- Swap space must be large enough to hold all processes

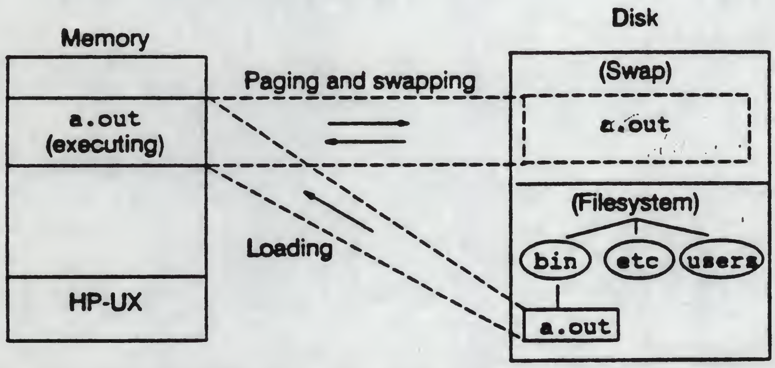


51436/51482 MS-4. 4 

## Student Notes

Physical memory is a finite resource on a computer. This means that only so many processes may fit into physical memory at any one moment in time, even though many more processes may actually be ready to run or execute. Because demand for physical memory may exceed actual supply, the concept of swap was introduced to computer systems. When demand for physical memory is high, entire processes or portions of processes are pushed out to the swap area on disk. When physical memory demand is low, entire process or portions of processes are brought back into memory from the swap area.

The **vhand** daemon monitors each page of available memory in order to discover pages that have not been recently accessed. If demand is high and pages exist that have not been recently accessed, **vhand** are pushes out or **pages out** these pages to the swap device. The **statdaemon** keeps track of how busy the **vhand** daemon is. As the **statdaemon** notices that **vhand** is unable to free up memory quickly enough



(99

to meet demands for memory, the `swapper` is called in to help. `vhand` pushes pages of a process out to the swap device; the `swapper swaps out` or pushes the entire process out to the swap device.

Most processes are shared text processes. A shared text process is a process that allows its code to be shared among all users currently executing its code. Each user, however, will have a copy of his own data, bss, and heap. The user's data, bss, and heap will have space allocated in the swap area in case of page or swap out. Since text is shared, it cannot be altered and therefore the original copy of the text from the executable file is always up to date. Swap space is not required for text, since we can always page in or swap in text from the executable file. A process' space on the swap device is reserved at process start up time; therefore, swap must be large enough to hold all executing processes.

## MS-5. SLIDE: Types of Swap Space

141

### Types of Swap Space

Device swap – a disk or section of a disk that is used exclusively for swap

- Primary swap - device swap which located on the system disk and is available at boot
- Secondary swap - device swap used in addition to primary swap

File system swap – a file system that also supports swapping

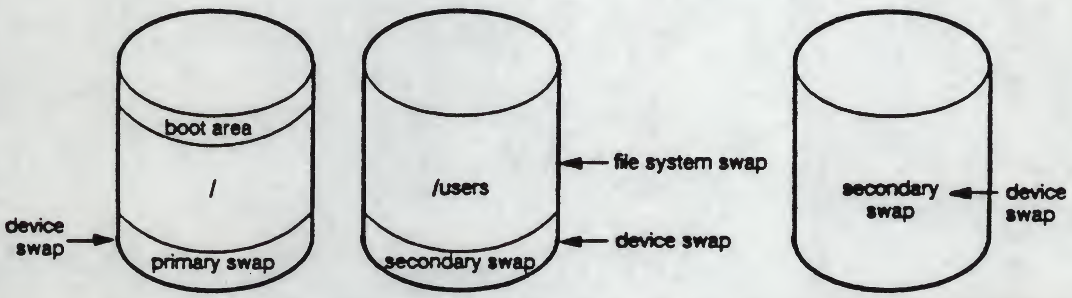


91436/91482 MS-5. 5 

## Student Notes

HP-UX has introduced some swap terms to better define its particular brand of swapping:

| | |
|---|---|
| Device swap | Device swap is a disk or section of a disk that is used exclusively for swap. |
| File system swap | File system swap, unlike the exclusively purposed device swap, is a file system that not only supports files and their data structures, but also has space available for swapping. |
| Primary swap | Primary swap is a special type of device swap which must be available at boot. Primary swap is located on the same disk as the root file system. |
| Secondary swap | Secondary swap is device swap that used in addition to primary swap and located on disk other than the root disk. |



200

Another term that can be seen from time to time is **dynamically allocatable swap**. This is disk space, device swap or file system swap, that can be allocated while the system is running. The only type of swap that is not dynamically allocatable swap is primary swap.

## MS-6. SLIDE: Evaluating Swap Space Needs

142

### Evaluating Swap Space Needs

When does a system administrator evaluate swap space needs?

- At System Installation Time
- If the following error message appears on the system console:

```
Sorry  pid  1622  was killed due to no swap space.
```

Swap space size estimation =
    sum of data and bss sizes for all processes
    + sizes of shared memory segments

91436/91482 MS-6. 6 

## Student Notes

A system administrator should always be monitoring system swap space, but there are times when this question becomes urgent to a system administrator. Before you install your system, you should try to get an idea of how much swap space you will need. Most application programs need a minimum amount of swap space to operate properly, and this figure is usually contained in the documentation that comes with the application.

Another event that may force you to re-evaluate swap space size is when the error message:

```
Sorry pid 1622 was killed due to no swap space.
```

appears on the system console indicating that processes are being terminated because there is no swap space available.

# Module MS — Memory and Swap Management

Estimation of swap size is not trivial, because swap space usage will vary with system load. Since what we are trying to estimate is maximun swap size needed, we need to understand swap size requirements when demand on the system is heaviest. In other words, at what moment in time will the programs that use the most data, bss, and shared memory all be running concurrently? (Bss is data that is uninitialized at the beginning of process execution.) We are interested in only data size, bss size and shared memory because these are all reserved swap space when a process is created. Swap space for text is not reserved. (A more generic term for text is a process' code.)

The best way to acquire information on process size is with the **size** command.

```
# size /usr/bin/vi
243332 + 22992 + 157620 = 423944
```

**size** returns the text, data, and bss sizes for a program. Combine the second and third sizes to estimate the size this program will occupy in the swap area. Don't forget that each user executing a process is allocated his own data and bss. For example, if there are 23 copies of `vi` executing at the system's busiest time, multiply `vi`'s data and bss by 23 ( (24332 + 22992) x 23 ) to account for all `vi`'s swap requirements. Repeat this calculation for each program on the system executing at the system's busiest moment.

Additionally, a process may also make use of shared memory. The `ipcs` command can be used to find out how much shared memory is in use at any one moment in time.

```
# ipcs -b
IPC status from /dev/kmem as of Tue Feb 26 18:11:10 1991
T      ID     KEY          MODE        OWNER     GROUP QBYTES
Message Queues:
q       0 0x3c201032 -Rrw--w--w-      root      root  16384
q       1 0x3e201032 --rw-r--r--      root      root    264
T      ID     KEY          MODE        OWNER     GROUP  SEGSZ
Shared Memory:
m       0 0x623120eb --rw-rw----      root       sys     80
T      ID     KEY          MODE        OWNER     GROUP NSEMS
Semaphores:
s       0 0x01090522 --ra-r--r--      root      root      1
```

The shared memory segment size is 80 in the above example.

Add the shared memory segment sizes to the data and bss sizes for an estimate. Add 256KB to this estimate in order to account for space used by the `exec()` system call during program initialization. Then to be conservative to account for page boundaries and the like, round up your estimate. Modify this estimate to be a multiple of 2MB. 2MB is the default size of **swchunk**, the kernel parameter for swap allocation unit.

## MS-7. SLIDE: Guidelines for Selecting Device Swap Areas

143

### Guidelines for Selecting Device Swap Areas

- Two swap areas on different disks are better than one single
- Series 600/800: Only one swap section per disk
- Device swap areas should be of similar size

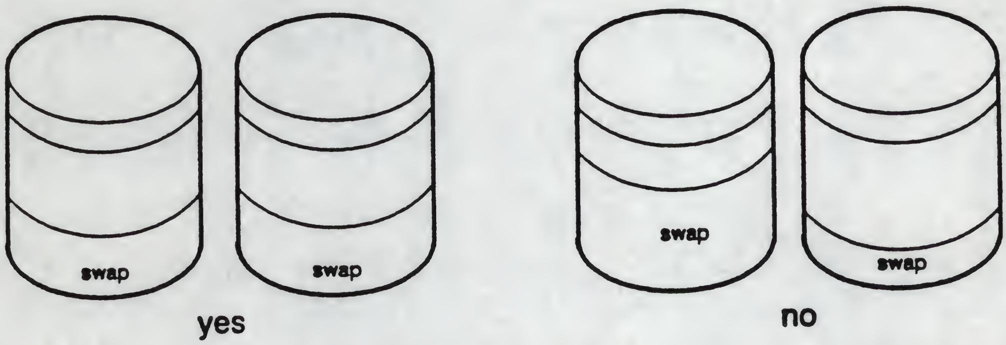


91436/91482 MS-7. 7 

## Student Notes

The guidelines for selecting swap areas are given above. Most of the recommendations are for performance reasons. You may set up swap space any way you like, but your system will run slower if the above rules are not applied.

From the point of view of performance, two swap areas on different disks are better than one swap area with the equivalent amount of space. Also for performance reasons, on the Series 600/800, multiple swap sections on the same disk should not be used. Multiple swap areas are used in an interleaved fashion. **Interleaved swapping** means that space from one swap device is used and then space from another space device. If the swap areas are on different disks, the swap areas can be written to concurrently avoiding disk head contention for multiple writes to the same disk.



202

**Device swap areas should be of similar sizes for best performance because, otherwise, when all space in the smaller device swap area is used, the larger swap area is all that is available and interleaving is no longer possible.**

## MS-8. SLIDE: Guidelines For Selecting File System Swap Areas

144

### Guidelines For Selecting File System Swap Areas

- Avoid using busy file systems such as the root file system
- Use `bdf` to check file system space availability
- Set priorities appropriately
  - Device swap over file system swap
  - Faster devices over slower devices
  - Seldomly used file systems over busier file systems

91436/91482 MS-8. 8 

## Student Notes

When you need more swap space and you have no disk space available for additional device swap, you can dynamically add file system swap to your system.

These guidelines are put forth to help the System Administrator decide on which file system sections should be used for swap space. Once again most of these rules stem from performance issues, as listed below.

When you add swap areas, you can assign a priority to each swap area. The system uses the swap areas with higher priority before using the lower priority swap areas. If you assign the same priority to two different swap areas, the system will use each of them on an alternating or round robin basis.



203

In general, it's best to assign highest priorities to the swap areas which afford the fastest performance. This means, give device swap areas priority over file system swap areas, give faster devices priority over slower devices, and give lower use file systems priority over higher use file systems. Actually, swap devices of the same type, e.g. all device swap on the approximately the same speed disk, should all be assigned the same priority in order to take advantage of interleaved swap.

## MS-9. SLIDE: Adding Swap Space Manually

145

### Adding Swap Space Manually

1) Add an entry to /etc/checklist

- Series 300/400/700:

```
block         directory  type    options
/dev/dsk/5s0  defaults   swap    defaults                        0 0
/dev/dsk/4s0  defaults   swap    end                             0 0
defaults      /users     swapfs  min=10,lim=4500,res=100,pri=0   0 0
```

- Series 600/800:

```
block              directory  type    options
/dev/dsk/c1d0s0    defaults   swap    defaults                        0 0
defaults           /users     swapfs  min=10,lim=4500,res=100,pri=0   0 0
```

2) Use swapon —a to enable entry

91436/91482 MS-9. 9 

## Student Notes

/etc/checklist is used by many different commands used by the system administrator. One of these such command is swapon -a. swapon -a searches /etc/checklist for entries with swap or swapfs in the third or *type* field.

If the *type* field is swap, *directory*, *backup-frequency*, and *pass_number* are ignored. Valid options are:

end (Series 300/400/700 only) This option is used when a device swap area shares the disk device with a file system. (The file system may or may not also be used for swapping. end ensures that the device swap area will start after the file system. Normally, if a file system exists on the device, swapon fails and outputs a warning message.



204

pri=*priority* — *priority* indicates the order in which space is taken from the file systems and devices used for swapping. Space is taken from the lower priority systems first. *priority* can have a value between 0-10. Default priority is 1.

If the *type* field is **swapfs**, *block*, *backup-frequency*, and *pass_number* are ignored. Valid options are:

min=*min* — *min* indicates the number of file system blocks the swap system will initially take from the file system. The default value for *min* is 0, indicating no swap space is to be allocated initially. Block size is defined by the administrator at the time the file system is created.

lim=*limit* — *limit* specifies the maximum number of blocks the swap system is allowed to take from the disk, provided space is available that is not reserved for exclusive use by the file system. The **res** option takes precedence over the **lim** option. The default value for *limit* is 0, indicating there is no limit to the amount of file system space the swap system can use. Block size is defined by the administrator at the time the file system is created.

res=*reserve* — *reserve* specifies the number of file system blocks in addition to the space currently occupied by the file system that are reserved for file system use only, making them unavailable to the swap system. The default value for reserve is 0 indicating that no file system space is reserved for file system use only. Block size is defined by the administrator at the time the file system is created.

pri=*priority* — *priority* indicates the order in which space is taken from the file systems and devices used for swapping. Space is taken from the lower priority systems first. *priority* can have a value between 0-10. Default priority is 1.

## MS-10. SLIDE: Using SAM to Add Device Swap (Series 300/400/700)

146

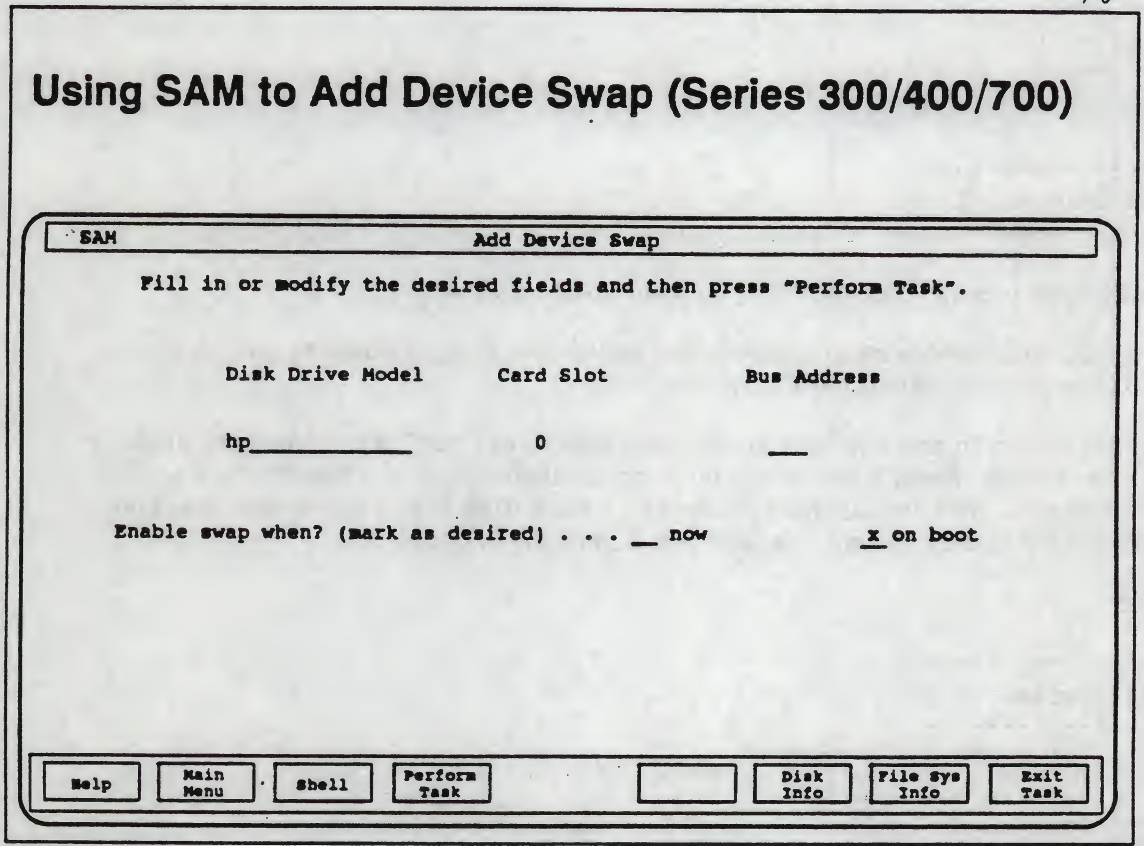


91436/91482 MS-10. 10 

## Student Notes

SAM allows you to:

- Add a disk to be used entirely for swap
- Add a disk to be used for file system and swap
- Remove device swap

SAM is the simplest way to add new device swap space.

To add swap, select:



205

```
--------------------------
| Peripheral Devices -> |
--------------------------
 |
 v
------------------
| Disk Drives -> |
------------------
 |
 v
-----------------------
| Add Device Swap ... |
-----------------------
```

Once in the Add Device Swap menu, select the disk you want and press Done.

Next select *when* you want your device swap enabled. You can choose now, on boot, or both. on boot will cause an entry to be created in /etc/checklist.

If you are not adding file system to this new disk at this time, select Perform Task. Otherwise, finish up in SAM by adding the file system. Specify the mount directory and answer as to whether you want the file system created now or not. With the last field on the Add a Hard Disk Drive menu, you can choose to view and/or modify the file system options. Answer n to accept the defaults.

To remove swap, select:

```
--------------------------
| Peripheral Devices -> |
--------------------------
 |
 v
------------------
| Disk Drives -> |
------------------
 |
 v
--------------------------
| Modify Device Swap ... |
--------------------------
```

Look at the device files that correspond to the disks containing swap areas; more the cursor to the entry for the devices you no longer want to use for swap. Enter n in the field labeled Enable on Boot.. Press Perform Task. The device swap area will be removed the next time you boot your system.

```
--------------------------
| Peripheral Devices -> |
--------------------------
 |
 v
------------------
| Disk Drives -> |
------------------
 |
 v
------------------------
| Add Device Swap ... |
------------------------
```

Once in the `Add Device Swap` menu, select the disk you want and press `Done`.

Next select *when* you want your device swap enabled. You can choose `now`, `on boot`, or both. `on boot` will cause an entry to be created in `/etc/checklist`.

If you are not adding file system to this new disk at this time, select `Perform Task`. Otherwise, finish up in SAM by adding the file system. Specify the mount directory and answer as to whether you want the file system created now or not. With the last field on the `Add a Hard Disk Drive` menu, you can choose to view and/or modify the file system options. Answer `n` to accept the defaults.

To remove swap, select:

```
--------------------------
| Peripheral Devices -> |
--------------------------
 |
 v
------------------
| Disk Drives -> |
------------------
 |
 v
---------------------------
| Modify Device Swap ... |
---------------------------
```

Look at the device files that correspond to the disks containing swap areas; more the cursor to the entry for the devices you no longer want to use for swap. Enter `n` in the field labeled `Enable on Boot.`. Press `Perform Task`. The device swap area will be removed the next time you boot your system.



206

## MS-11. SLIDE: Using SAM to Add File System Swap

147

### Using SAM to Add File System Swap

```
SAM                          Add File System Swap

File system swap is taken from the file system and should be used to augment
other swap rather than replace it. File system swap should not fill up the
file system. Move the cursor to line of the file system you wish to change,
modify the desired fields and then press "Perform Task".

Mount Directory     Minimum Swap     Maximum Swap     Total File System
                      (kbytes)         (kbytes)        Size (kbytes)

/                     0                0                 128968
/usr                  0                0                 161578

 Help   Main Menu   Shell   Perform Task        Disk Info   File Sys Info   Exit Task
```

91436/91482 MS-11. 11 

## Student Notes

Using SAM to Add File System Swap SAM allows you to:

- Add a file system for swapping
- Change current values for `minimum` and `maximum` swap space

Once again SAM is the simplest way to add new file system swap space.

To add file system swap, select:

```
-------------------------
| Peripheral Devices -> |
-------------------------
```

```
 |
 v
------------------
| Disk Drives -> |
------------------
 |
 v
----------------------------
| Add File System Swap ... |
----------------------------
```

Once in the Add File System Swap menu, SAM lists the file systems on your system for use as swap. Select the desired file system. You can also select File Sys Info. This screen will show you the space available in all of the mounted file systems.

Specify value for Minimum Swap for minimum amount of the file system to use for swap. Specify a file system swap space limit with Maximum Swap.

Next SAM will ask:

`Do you wish to enable the file system swap now? (y or n)`

which will enable swapping, as the question says, immediately.

SAM follows with:

```
Do you wish to enable the file system swap permanently?  (only permanent
file system swap appears on this screen, but temporary file system swap is
visible on the view disk information screen) (y or n)
```

If you answer y, SAM will make an entry in /etc/checklist.

SAM help screens provide instructions for changing the values currently set for minimum and maximum swap space.



207

## MS-12. WORKSESSION: Review Questions

1. Define physical memory and virtual memory.

2. What is the difference between swapping and paging?

3. Define physical memory, available memory, and lockable memory.

4. What are the two types of swap space that a system administrator can set up?

5. What are the the two types of device swap?

6. What are the advantages of device swap?

7. What are the advantages of file system swap?

8. Name the two methods for allocating swap space?

9. Describe two guidelines for setting up device swap.

10. Describe guidelines for setting up dynamic swap space.

11. Describe how to set priorities for dynamic swap.

12. What is the name of the file used to automatically setup either device or file system swap?



208

## MS-13. LAB: Memory and Swap Management (Series 300/400/700)

### Directions

Your instructor will advise you as to what disk device to use for this lab. Make sure you use this disk for this lab.

1. How much memory does your training system have?

2. How much of the memory is lockable?

3. Manually add your disk as a dedicated swap disk (no file system area).

4. Remove the disk as a dedicated swap disk.

On the disk, create and mount a file system and create device swap.

5. Enable file system swap on the disk.

## MS-14. LAB: Memory and Swap Management (Series 600/800)

### Directions

**Your instructor will advise you as to what disk section to use for this lab. Make sure you use this disk section for this lab.**

1. **How much memory does your training system have?**

2. **How much of the memory is lockable?**

3. **Manually add your disk section as swap.**

4. **Remove the swap section you just created.**

**On the disk, create and mount a file system.**

**Enable file system swap on the disk.**

# Solutions

## MS-16. LAB: Memory and Swap Management (Series 600/800)

1. Define physical memory and virtual memory.

**Answer:**

**Virtual address space** is all valid memory addresses; whereas **physical address space** is the actual RAM purchased.

2. What is the difference between swapping and paging?

**Answer:**

**Swapping** involves moving entire processes between RAM and mass storage (usually disk), whereas **paging** involves moving smaller units (called pages) between RAM and mass storage.

3. Define physical memory, available memory, and lockable memory.

**Answer:**

**Physical memory** is the actual RAM purchased. That physical memory is divided up for different uses by the MMU. Some memory is reserved for kernel code and data structures. The amount of memory remaining is referred to as **available memory**, and is used by the system for demand paging. All or part of available memory can be locked by a subsystem or by user processes. **Lockable memory** cannot be swapped out to disk. Typically, locked memory holds frequently accessed programs or data structures.

4. What are the two types of swap space that a system administrator can set up?

**Answer:**

Device swap and file system swap.

5. What are the the two types of device swap?

**Answer:**

Primary and secondary swap.

6. What are the advantages of device swap?

**Answer:**

Faster.

7. What are the advantages of file system swap?

**Answer:**

Less disc space totally committed to swap space.

8. Name the two methods for allocating swap space?

**Answer:**

Manual method, or SAM.

9. Describe two guidelines for setting up device swap.

**Answer:**

1. Two swap areas on different discs better than one larger swap space.
2. Swap areas on different disc should be of the same size.

10. Describe guidelines for setting up dynamic swap space.

**Answer:**

1. Avoid using heavily used file systems, like the root file system.
2. Set priorities wisely.

11. Describe how to set priorities for dynamic swap.

**Answer:**

Select device swap over file system swap.

Select faster devices over slower devices.

Select lower used file systems over highly used file systems.

12. What is the name of the file used to automatically setup either device or file system swap?

**Answer:**

`/etc/checklist`.

1. How much memory does your training system have?

**Answer:**

Use the `dmesg` command to see `real mem`, `avail mem`, and `lockable mem`.

2. How much of the memory is lockable?

**Answer:**

Varies depending on system.

3. Manually add your disk as a dedicated swap disk (no file system area).



2.10

# Solutions

**Answer:**

1. Add the something like the following to `/etc/checklist`:

```
/dev/dsk/3s0 defaults   swap   defaults  0 0
```

2. Execute `swapon -a` or reboot to enable this swap area.

4. Remove the disk as a dedicated swap disk.

On the disk, create and mount a file system and create device swap.

**Answer:**

Again using SAM or the `/etc/swapon` command.

1. For the file system:

```
# newfs /dev/rdsk/3s0 hp2213A
# mount /dev/dsk/3s0 /aux
```

2. For the device swap:

   Add the something like the following to `/etc/checklist`:

```
/dev/dsk/3s0 defaults   swap   end                         0 0
```

   Execute `swapon -a` or reboot to enable this swap area.

5. Enable file system swap on the disk.

**Answer:**

1. Add the something like the following to `/etc/checklist`:

```
defaults    /aux    swapfs min=10,lim=4500,res=100,pri=0    0 0
```

2. Execute `swapon -a` or reboot to enable this swap area.

1. How much memory does your training system have?

**Answer:**

Use the `dmesg` command to see `real mem`, `avail mem`, and `lockable mem`.

2. How much of the memory is lockable?

**Answer:**

Varies depending on system.

3. Manually add your disk section as swap.

# Module RK — Reconfigure Kernel (s300/400/700)

## Objectives

Upon completion of this module, you will be able to:

- Describe reasons why it might be necessary to reconfigure the kernel.
- Perform the steps to reconfigure the kernel.
- Reconfigure the kernel for an additional device or subsystem.
- Describe some important tunable parameters.



215

# Module RK — Reconfigure Kernel (s300/400/700)

## RK-1. SLIDE: Overview of Reconfiguration

1 98

### Overview of Reconfiguration

Reasons for reconfiguring the kernel include:

- Adding or deleting devices
- Adding or deleting subsystems
- Changing tunable parameters

Reconfigure interactively using manual commands or using SAM

51436 RK-1. 1 

### Student Notes

For most systems, the default kernel configuration supplied by Hewlett-Packard is sufficient and you will not have to modify the HP-UX kernel configuration. However, if you do need to change the default configuration, you will have to modify the kernel. Reasons for reconfiguring the kernel include:

- Adding or deleting devices
- Adding or deleting subsystems
- Changing tunable parameters

You can modify the kernel either by using SAM, the System Administration Manager utility, or by manually entering commands.

# Module RK — Reconfigure Kernel (s300/400/700)

## RK-2. SLIDE: The Reconfiguration Process

149

### The Reconfiguration Process

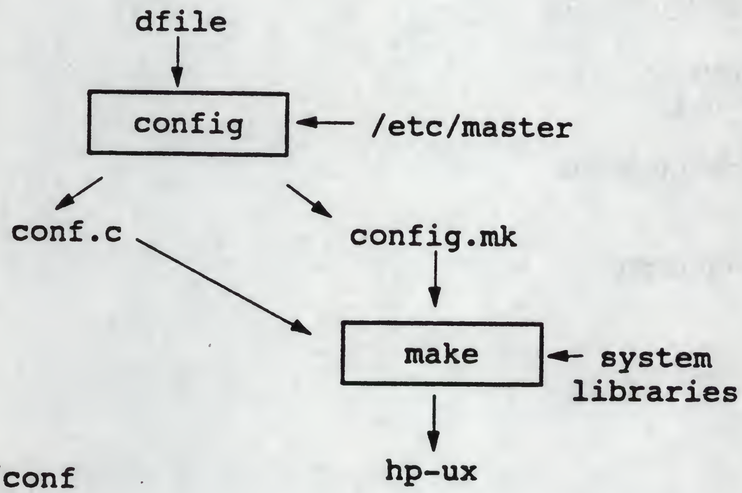


```
#cd /etc/conf
#config dfile
#make -f config.mk
```

51436 RK-2. 2 

## Student Notes

The reconfiguration process starts with the `dfile`. The `dfile` is the system description file. It specifies what drivers are used in the kernel, and optionally specifies root and values for system parameters. `/etc/master` and `dfile` are the input to `config`. `/etc/master` contains all the information regarding supported devices. `/etc/master` is supplied as part of the HP-UX operating system and should not be modified by anyone who does not fully understand its structure and purpose.

`config` generates two output files. One of these files, `conf.c` is a C program that defines the configuration table for the various devices on the system. The other output file is `config.mk`. `config.mk` holds instructions on how to create the HP-UX kernel. `make`, using `config.mk`'s instructions, compiles and links `conf.c` and outputs `hp-ux`, the new HP-UX kernel. `make` is a standard hp-ux utility that can be simply described as a program construction tool.

## RK-3. SLIDE: The `dfile`

150



### Student Notes

A **`dfile`** contains two parts. The required part of the **`dfile`** includes:

- device drivers
- pseudo-drivers, such as ptys

Each line has the following format:

*devname*

where *devname* is the driver name for the device or interface card as it appears in the alias table in the file `/etc/master`. For example, `scsi` selects the SCSI disk drives. `/etc/master` contains the complete list of configurable devices, cards, and pseudo-drivers.

Typically after installation only a minimum number of drivers are installed and so you will want to add other drivers. Your system will work just fine if you have drivers configured into your kernel that are not actually on your system. In other words, you can have entries in your `dfile` for devices that are not really on your system. Remember though, that your kernel will be bigger and will take up more memory that cannot be reclaimed.

Only one entry is needed in your `dfile` for each category of devices. For example, if you have several `cs80` devices on your system, you only need one `cs80` entry in your `dfile` and if you have several SCSI devices on your system, you only need one `scsi` entry in your `dfile`.

On a series 700 system, the following devices are not configurable and should not be specified in the `dfile`:

```
cons hil hilkbd ite    framebuf
graph3 mm nm netdiagl   r8042
tty
```

On a series 300/400 system, the following devices are not configurable and should not be specified in the `dfile`:

```
cons  hil ite200 nimitz sy
graphics iomap mm r8042 tty
```

The optional part of the `dfile` includes:

root — This is usually not explicitly defined in the `dfile`. HP-UX will assume the root disk is the same disk that was used initially at boot up.

swap — For a system with less than eleven device swap areas, this entry does not need to be present.

system parameters — These parameters are usually left as installed unless there is a compelling reason to change them such as an application program exhausting a certain system resource.

For system parameters, each line contains two fields:

`parameter_name` *number or formula*

Categories for these parameters are:

Message Related Parameters
Semaphore Related Parameters
Shared Memory Related Parameters
Accounting Code Parameters
File System Parameters
Process Parameters
Miscellaneous Parameters
Networking Parameters

# Module RK — Reconfigure Kernel (s300/400/700)

Cluster Related Parameters
Hardware Related Parameters
MS-DOS Related Parameters (optional for machines with DOS co-processors)

These parameters are documented in detail in appendix A of the *HP-UX System Administration Tasks* manual.

## RK-4. SLIDE: Sample dfiles

151

### Sample dfiles

**Table 7-1.**

| Sample | Description | Function |
|---|---|---|
| `dfile.min` | Contains one of the minimal operating system configurations. | Minimal O.S. on a standalone system |
| `dfile.full` | Configures HP-UX with all the supported device drivers in this release, minus LAN, NFS, cluster drivers, and pseudo drivers other than `pty`. | A fully loaded system less LAN |
| `dfile.full.lan` | `dfile.full` plus the LAN driver. | Default (and fully loaded) O.S. for a standalone system |
| `dfile.cnode` | Contains a minimum kernel configuration for a client on an HP-UX cluster. | Default (and minimal) O.S. for client |
| `dfile.cnodemax` | Same as `dfile.maxservr` except replace cluster server-specific information with client- specific information. | Fully loaded O.S. for client |
| `dfile.cnodemin` | The minimum kernel configuration for a client (will not work on many clients); included only as a sample of a bare minimum kernel. | Illustration |
| `dfile.maxservr` | Contains the drivers in `dfile.full.lan`, plus the cluster drivers and operating system parameters required for a cluster server. | Default (and fully loaded) O.S. for cluster server |
| `dfile.minservr` | Contains all drivers in `dfile.min`, plus the cluster drivers and operating system parameters required for a cluster server. | Minimal O.S. for a cluster server |

51436 RK-4. 4 

## Student Notes

There are several samples of a `dfile` in `/etc/conf`. If you choose one of these samples, you should copy the sample to the current `dfile`. It is wise to not to make changes to the original copy of these samples. Also, backup the existing `dfile` before modifying it. The table describes each of these samples.



2/8

## RK-5. TEXT PAGE: dfile.full.lan for Series 700

```
* Installed software drivers and I/O interface cards
* Magneto-Optical autochanger
autox
autoch
* SCSI disk
scsi
* SCSI DAT
scsitape
* SCSI-2
sim0
* parallel interface driver
parallel
* Pseudo-entries to bring in the 4.3-based networking initialization
*   Network Management
netman
* SLIP driver
ni
*   NS networking
nipc
*   domain sockets - needed for X
uipc
*   internet domain sockets
inet
* Networking drivers
lan01
asio0
* Network File System (NFS)
nfs
* CD File System (CDFS)
cdfs
* Pty pseudo drivers
pty0
pty1
* Kernel test drivers
*ktest
*ktestio
* libIO Kernel interface
dconfig
* I/O Diagnostics
diag1
* Memory Diagnostics
dmem
* Diskless cluster code pseudo driver
*dskless
* Remote Swap for DUX
*rdu
* Apple-Talk
*atalk
```

```
*
*xns
eisa
CharDrv
*
* Tunable parameters
*
*dskless_node 1
*server_node 1
* Swap info
* swap auto     /* Swap on root disk. */
* swap scsi E0100 -1 /* Use only the swap partition of the disk. */
```

## RK-6. TEXT PAGE: dfile.full.lan for Series 300/400

```
* Installed software drivers and I/O interface cards:
*
* Most mass storage devices (exists in all config. files)
cs80
* SCSI direct access storage devices
scsi
* Amigo mass storage device driver
amigo
* 9-track magnetic tape drives
tape
* 9-track streaming tape drives
stape
* Ciper printer
ciper
* Non-ciper printer
printer
* Shared Resource Manager driver
srm
* Pseudo terminal drivers (required for X windows, etc.)
ptymas
ptyslv
* HP-IB driver; also include for Device I/O Library (DIL)
hpib
* GPIO driver; also include for DIL
gpio
* Remote Job Execution
rje
* HP 98286 DOS Coprocessor driver
dos
* HP 98646 VME driver
vme
* HP 98577 VME expander
vme2
* Bi-directional parallel interface driver
parallel
* Network File System
nfs
* Pseudo-entries for BSD 4.3-based networking initialization
uipc
nipc
inet
netman
* CD File System (CDFS)
cdfs
* Measurement System driver
meas_sys
* Kernel Instruction Coverage Analyzer
*ica
```

```
* Networking drivers
lla
lan01
netdiag1
* nsdiag0
* X.25 driver info
x25
x25ip
x25pa
pdix25
*
* CARDS:
*
* HP-IB interface
98624
* high-speed HP-IB interface
98625
* RS-232 serial interface
98626
* RS-232 datacom interface
98628
* RS-232 multiplexer
98642
* Internal RS-232 serial interface
apci
* SCSI interface
98265
*
* Tunable parameters:
*
* Swap info:
*
```



216

## RK-7. SLIDE: Steps to Manually Reconfigure the Kernel

152

### Steps to Manually Reconfigure the Kernel

1. Bring your system to a single-user state
2. `#cd /etc/conf`
3. Modify the `dfile` to make changes to kernel configuration
4. `#config dfile`
5. Verify that `config` created `conf.c` and `config.mk`
6. `#make -f config.mk`
7. Verify that `make` created `hp-ux` in `/etc/conf`
8. `#mv /hp-ux /SYSBCKUP`
9. `#mv hp-ux /hp-ux`
10. Reboot the system

51436 RK-7. 5 

### Student Notes

Before you begin the process of reconfiguring the kernel, bring your system to single-user mode and then change working directory to the directory `/etc/conf`.

```
# cd /
# shutdown 0
  .
  . wait for shutdown to complete
  .
# cd /etc/conf
```

The commands to perform the steps on the slide are shown below:

- Modify `dfile` to include your changes:

  The `dfile` is a description of an HP-UX system configuration. If possible, make modifications to your existing `dfile`. You might want to make a backup copy first in case you want the original again.

  ```
  # cp dfile dfile.old
  # vi dfile
  ```

  If major modifications are required to the current `dfile`, check the samples in `/etc/conf` to see if they might be closer to what you want (for example, `dfile.full` or `dfile.min`). If you choose to use one of these samples, you should copy the sample to the current `dfile`.

  ```
  # cp dfile.sample dfile
  # vi dfile
  ```

- To use the `config` command to process `dfile`:

  ```
  # config dfile
  ```

- Verify that `config` created `conf.c` and `config.mk` in your local directory

  ```
  # ls
  ```

- Run `make` on `config.mk` to regenerate the kernel:

  ```
  # make -f config.mk
  ```

- Verify that `make` created `hp-ux` in your current directory:

  ```
  # ls
  ```

- Back up the existing kernel. Remember, you are in the directory `/etc/conf` so full path names are required for this step:

  ```
  # cp /hp-ux /SYSBCKUP
  ```

---

**Note** Do NOT do this on a cluster client. Though this is not a class on the HP-UX Clustered Environment, this note is for future reference, since some System Administrators may have such an environment.

---

- Move the new kernel to `/hp-ux` and reboot the system:

  ```
  # mv hp-ux /hp-ux
  # /etc/reboot
  ```

# Module RK — Reconfigure Kernel (s300/400/700)

**Note** For an HP-UX cluster, you must reconfigure the kernel from the cluster server. You cannot use a cnode except to rebuild the kernel for that client. See comment about NOT backing up the kernel.

## If the New Kernel Won't Boot, Boot /SYSCKUP

### Series 700

To boot the backup kernel, /SYSBCKUP, redirect the boot sequence provided by the Boot Console User Interface by hitting the ESCAPE key until the message,

`Terminating selection process.`

is displayed.

The automatic boot sequence has now been halted and you are in fully "attended" or interactive mode.

```
32 MB of memory configured and tested.

Selecting a system to boot.
To stop selection process, press and hold the ESCAPE key.

Selection process stopped.

Searching for Potential Boot Devices.
To terminate search, press and hold the ESCAPE key.

Device Selection     Device Path              Device Type
-----------------------------------------------------------------------

P0                   scsi.6.0                 QUANTUM PD210S
P1                   scsi.5.0                 QUANTUM PD210S

b)     Boot from specified device
s)     Search for bootable devices
a)     Enter Boot Administration mode
x)     Exit and continue boot sequence
?)     Help

Select from menu:  b p0 ipl

ISL> hpux boot disk(scsi.6;0)/SYSBCKUP
```

### Series 300/400

Boot in attended mode (that is, press the space bar right away during boot up). Booting in attended mode allows you to choose the operating system that you want to boot. At the menu, choose SYSBCKUP:

```
:HP7935 1400, 0, 0
```

```
1H SYSHPUX
1B SYSBCKUP
```

You will boot your old kernel, the one you copied into /SYSBCKUP by typing "1B" in this example to select the backup kernel.



218

## RK-8. SLIDE: Adding I/O Not Already in the Kernel

153

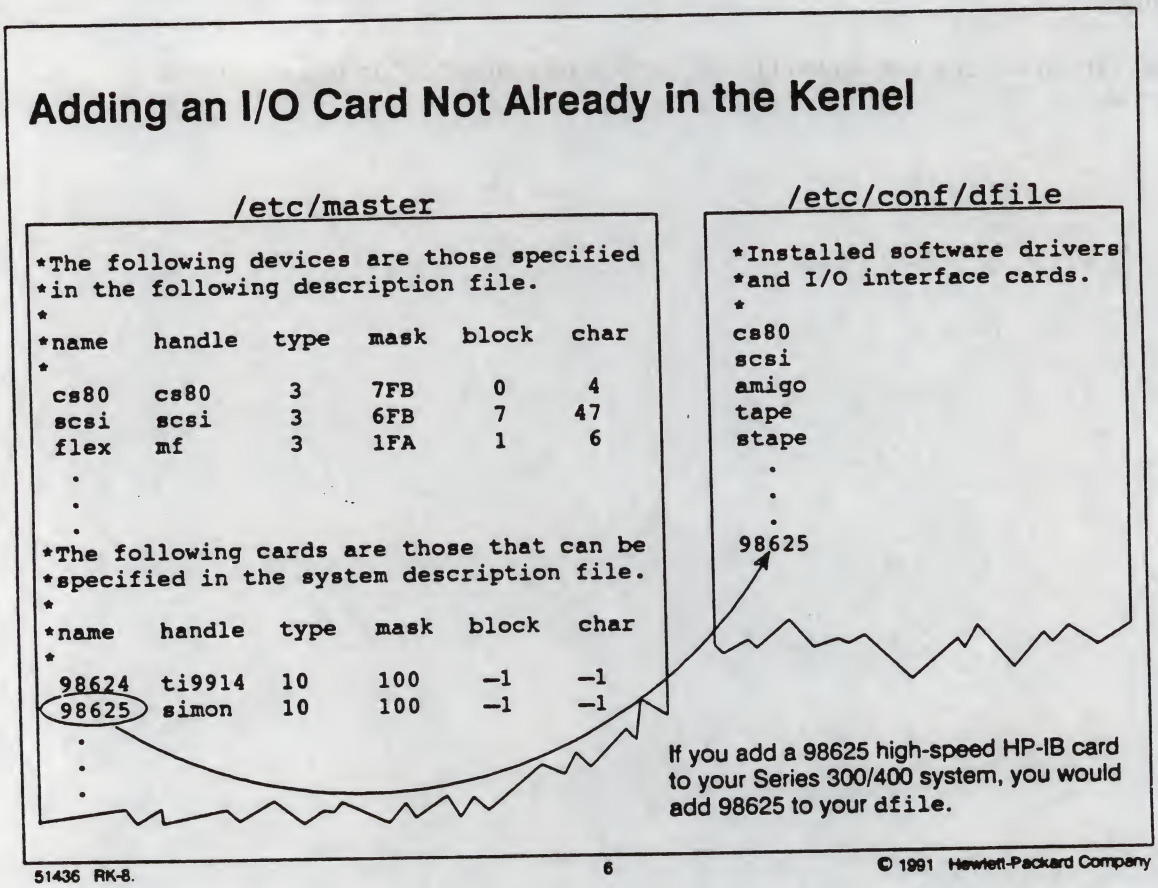


### Student Notes

### Student Notes

The **dfile** contains an entry for all peripheral devices, I/O cards, and major software subsystems such as LAN, RJE, X25, etc. To add a device, card, or subsystem to the kernel, you must add an entry for the device, card, or subsystem to your **dfile**. To determine what entry to add to the **dfile** you should examine the **/etc/master** file. The **/etc/master** file is the device information table that contains the device name, handler name, element characteristics, functions for the device, major device number, and minor device number. (We will look at this file in the module on Creating Device Files.) If what you are trying to add is not in the **/etc/master** file then it is not supported by Hewlett-Packard.

The example on the slide assumes you want to add a 98625 High Speed HP-IB card to your Series 300/400 system. To configure your kernel for this card, you would check the **/etc/master** to make sure it's supported and then add 98625 to your **dfile**. (You could also add **simon** to the dfile. Either the "name"

```
/etc/mknod /dev/dsk/0s0 b 7 0x0e0600
/etc/mknod /dev/rdsk/0s0 c 47 0x0e0600
```

2

```
open("/dev/dsk/0s0",O_RDWR)
```

1

drivers

interfaces

devices

```
/etc/master

*
* name          handle        type   mask   block   char
*
cs80            cs80          3      7FB    0       4
scsi            scsi          3      6FB    7       47

*
* name   handle  type   mask   block   char
*
98265    scsi    10     100    -1      -1

*
* The following entries form the alias table.
* field 1: product #    field 2: driver name
*
7958S   scsi
7959S   scsi
```

/etc/lsdev

dev afwezig: ENXIO

/etc/conf/dfile → /etc/master → conf.c ⟶ make -f config.mk ⟶ hp-ux

3

scsi
of
7958S

```
struct bdevsw bdevsw[] = {
/* 7*/  scsi_open, scsi_close, scsi_strategy, scsi_dump, scsi_size, C_ALLCLOSES, };

struct cdevsw cdevsw[] = {
/*47*/  scsi_open, scsi_close, scsi_read, scsi_write, scsi_ioctl, seltrue, C_ALLCLOSES, };
```

or the "handle" will work in the dfile.) Then follow the steps to configure your kernel manually as shown earlier in this module. They are listed again here for reference.

```
# vi dfile
  .
  .  edit dfile as required to add 98625
  .
# config dfile
# make -f config.mk
# mv /hp-ux /SYSBCKUP    (if this applies to your system)
# mv hp-ux /hp-ux

# cd /
# shutdown -h
```

- When "halted" message appears, power off the system.
- Insert interface card while following precautions about ESD protection.
- Power up your system.



220

## RK-9. SLIDE: Adding a Subsystem to the dfile

154

### Adding a Subsystem to the `dfile`:

- Procedures vary depending on the subsytem
- Often requires more changes than simply a modification to `dfile`
- Usually done with `/etc/update`

Example of `dfile` modifications for CD-ROM subsystem:

```
* CD File System (CDFS)
cdfs
```

51436 RK-9. 7 

### Student Notes

In almost all cases, you should use `/etc/update` to add a subsystem, rather than the manual kernel reconfiguration. Manual kernel reconfiguration will add the subsystem "driver", but will not get the software onto your system or do any other set-up that may be needed. `update` on the other hand, helps you load the software and does the necessary configuration.

When adding any subsystem, the documentation for that subsystem needs to be consulted. The subsystem documentation, not the documentation that applies to the system management in general, details what lines need to be added to the `dfile` and other changes to system such as the addition of software and device files.

This is an example of what changes are needed for the `dfile` in order to access a CD-ROM. This is *not* all the modifications necessary for CD-ROM to function on the system.



221

## RK-10. SLIDE: Reconfiguring the Kernel with SAM

155

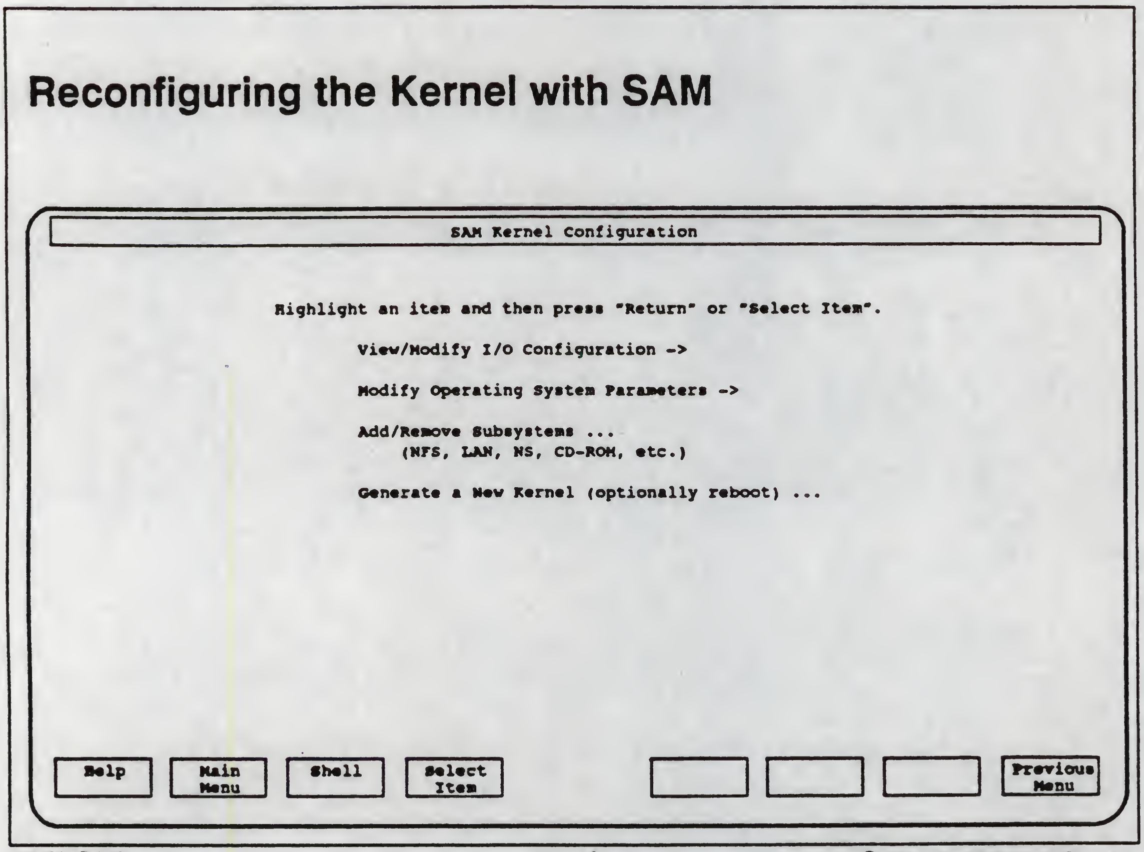


51436 RK-10. 8 

### Student Notes

This slide shows you the kernel reconfiguration tasks that SAM can perform for you. As you can see, SAM is quite versatile in this area. In most cases where you need to reconfigure the kernel, you will probably want to use SAM.

To get to the menu shown on the slide, choose `Kernel Configuration ->` from SAM's Main Menu. Before SAM continues, you will get a message that says:

```
The file below provides current kernel values for viewing or
modification.  To use a different set of initial values,
specify your own file.  In either case, press "Return" or
"Done" to proceed.
```

# Module RK — Reconfigure Kernel (s300/400/700)

```
File: /hp-ux
```

The file `/hp-ux` contains your executable kernel. This is the file that SAM will rebuild. (This is also the file that is rebuilt when you execute `make -f config.mk` manually.) If you do not want SAM to use the kernel values in `/hp-ux`, you can enter a different file name at the prompt. For example, if you just installed HP-UX 8.0 and this is the first time you are rebuilding your kernel, you might want to use the kernel configuration in the file `/etc/conf/dfile.full.lan`.

After you press (Return) you will see a brief message:

```
Initializing for Kernel Configuration
```

## RK-11. SLIDE: Adding a Subsystem with SAM

156

### Adding a Subsystem with SAM

```
SAM Add/Remove Subsystems

Select subsystems to add to or remove from the current configuration
and then press "Perform Task".

                              In configuration?
                                 (y or n)
        NS/9000 . . . . . . . . . . . . . y
        LAN/9000 . . . . . . . . . . . . y
        NFS/9000 . . . . . . . . . . . . y
        CD-ROM/9000 . . . . . . . . . . y
        DSKLESS/9000 . . . . . . . . . . n
        X.25/9000 IP Access . . . . . . n
        X.25/9000 Programmatic Access . n

Help   Main Menu   Shell   Perform Task                  Exit Task
```

51436 RK-11. 9 

## Student Notes

To get to this menu, select **Add/Remove Subsystems ...** from the Kernel Configuration Menu.

With the arrow keys, move to the subsystem you want to add. Type y over the n and then press **Perform Task**.

You will get a message similar to the following:

```
Subsystem information has been modified.  You will be
prompted to generate a new kernel to effect changes.

   -- Press the space bar to continue. --
```

Press the (Space) bar and then press `Main Menu` which will tell SAM you have made all the changes you want. You will get this message:

```
Kernel configuration information has been modified.  To
preserve the new kernel configuration, you must regenerate
the kernel.


Mark a choice with an "x" and press "Done".


x Regenerate the kernel with modified values.
_ Exit kernel configuration canceling changes.
_ Continue configuring the kernel.
```

Choose the first option and then press `Done`. SAM will create a new kernel with the changes you specified.

## RK-12. SLIDE: Changing System Parameters With SAM

157

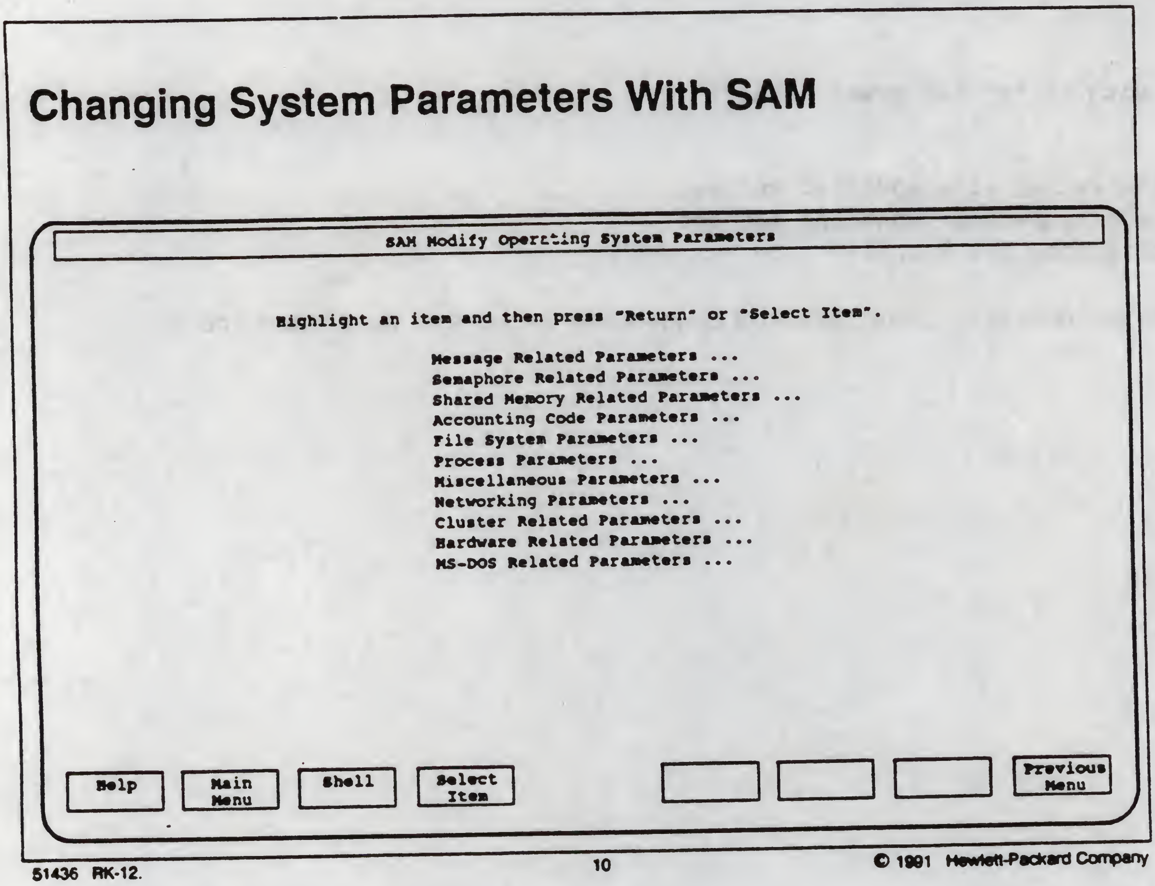


### Student Notes

To get to this menu, choose "Modify Operating System Parameters" from the Kernel Configuration menu. All of the operating system parameters are divided into groups. These groups are shown by SAM (see slide). To change a particular parameter, you must know what group it is in. See Appendix A of *HP-UX System Administration Tasks* for a complete listing of all the parameters. Once you know what group your parameter is in, highlight it and press `Select Item`.

If you were to choose "Process Parameters ... ," you would get the following new menu:

```
SAM                        Process Parameters

      Fill in or modify the desired fields and then press "Perform Task".

     Maximum number of processes which may
           simultaneously exist (nproc) . . . . . . . . . . . . .  84

     Maximum number of processes that a
           user may have (maxuprc) . . . . . . . . . . . . . . . . 50

     Maximum process data segment size in bytes (maxdsix) . . .   16777216

     Maximum process stack size in bytes (maxssiz) . . . . . . .  2097152

     Maximum process shared text segment
           size in bytes (maxtsiz) . . . . . . . . . . . . . . . . 16777216
```

Use the arrow keys to get to the parameter you wish to change and then type in the new number. Press `Perform Task` when you are finished.

---

**Note**

You should never modify operating system parameters unless you know exactly what will happen when you make a particular change. This is an advanced administrative task and should not be taken lightly. You can render your system unusable by improper modifications to kernel parameters

---

## RK-13. LAB: Reconfiguring the Kernel

### Directions

Working with your lab partner, you will reconfigure the kernel on your lab system. Perform the tasks described below.

1. Move to the `/etc/conf` directory. Copy the current `dfile` to `dfile.`*your_name*.

2. Now, edit `dfile.`*your_name* and add the appropriate entry for a CD-ROM file system if one is not already present. Also you should make sure your system can access the CD-ROM disk interface. CS-ROMs are supported on high speed HP-IB and SCSI on Series 300/400. The SCSI interface is supported on the Series 700.

3. Now, perform the manual steps to add these new device to your kernel. Once again, assume the hardware has already been added to the system.

4. Invoke SAM. Go to the "Kernel Configuration" menu. Look through SAM's menus on kernel configuration, but don't reconfigure your kernel with SAM.

5. Series 300/400: Invoke SAM. Go to the "Kernel Configuration" menu. Use SAM's menus on kernel configuration to reconfigure your kernel so that your system can talk to disks that support the AMIGO protocol.

6. Series 700: Add the appropriate entry for the parallel interface.



225

# Module UH — Updating HP-UX

## UH-1. SLIDE: Installing Versus Updating HP-UX

158

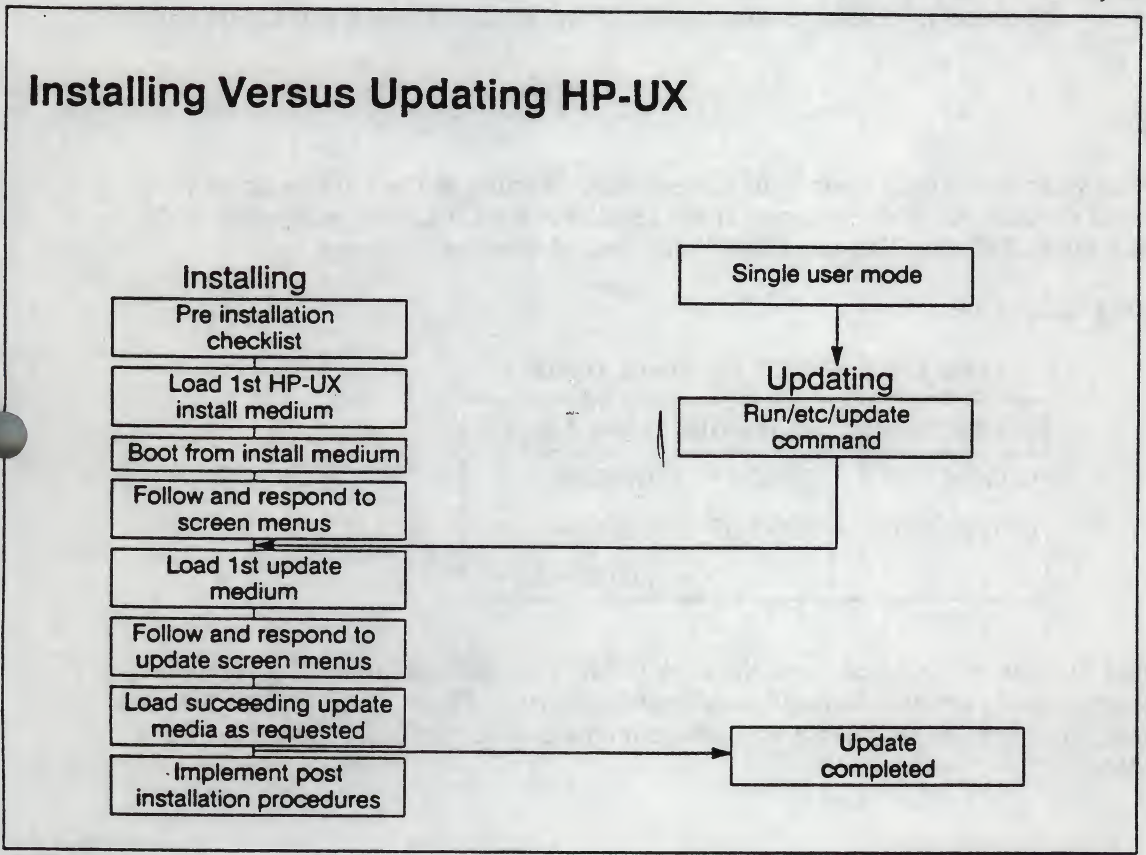


51436 UH-1. 1 

## Student Notes

The update process is the same as the "update" phase of the "install" process we covered in an earlier module. The relationship between the two processes is shown on the slide.

As a quick review of what we covered in the installation module: The install utility initialized the system disk, created a new file system, and copied over core HP-UX partitions and filesets and then rebooted the system. Then the install utility ran the **/etc/update** program to load the rest of the partitions and filesets to the system disk.

Remember that a **fileset** is a logical grouping of files and a **partition** is a logical grouping of filesets. A **product**, for example, the X Window System, is a logical grouping of filesets and/or partitions. When you load the filesets onto the disk from the update media, the files are copied to the disk in a hierarchical structure that is unrelated to the original fileset groupings on the update media.

When Hewlett-Packard releases a new version of the HP-UX operating system, such as version 8.0 or 9.0, you will use the update process to replace outdated portions of HP-UX with the current versions and add



226

new features. When you update HP-UX, you load software products from the update media (or from a special network server) and incorporate that software (in the form of files and filesets) into an existing HP-UX file system. If any new or replacement files affect the existing kernel configuration, the kernel is reconstructed during the update. If the kernel is reconstructed, the update process will reboot your system.

## Update Media Types

The type of update media you use depends upon your system type. Starting at the 8.0 release, new media types are supported on most HP 9000 systems. There are differences regarding supported update media types between the series 300/400, 700, and 600/800 families of computer systems.

Possible choices for media options are:

**Table UH-1. HP-UX 8.0 Media Options**

| Series 300/400 | Series 700 | Series 600/800 |
|---|---|---|
| CDROM | CDROM | DDS/DAT |
| cartridge tape | DDS/DAT | mag tape |
| | | cartridge tape |

You may also update your system over a Local Area Network (LAN) provided that there is a computer system that has been configured as a network distribution (netdist) server. This method will be described in the end of this module. First we will talk about the update methods that use local physical media instead of network updates.

## UH-2. SLIDE: Before You Update HP-UX

159

### Before You Update HP-UX

- Read the *HP-UX Release Notes* for the new release.
- Obtain codewords and hardware IDs if using codeword protected media.
- Make a note of the device file name for your installation device.
- Ensure that your `TERM` environment variable is set correctly.

51436 UH-2. 2 

## Student Notes

### Read the *HP-UX Release Notes* for the new release.

This slide shows the steps you should take before you perform your update. You should carefully follow these steps, modifying them as directed by other manuals needed when you situation differs from a standard system (particularly if updating an HP-UX Clustered Environment).

Prior to updating the operating system, you should read the *HP-UX Release Notes*. Any new release will include numerous changes that may affect your users and applications running on your system. The *HP-UX Release Notes* is available in electronic form in the file `/etc/newconfig/8.0ReleaseNote` after you have completed an installation of the HP-UX 8.0 release. The *README FIRST* document describes how you can extract this file from your **UPDATE** media before you start the update process.

The *Software Status Bulletin* (SSB) provides information about currently outstanding problems and enhancement requests. You can get information, for example, about a known problem

related to an HP-UX subsystem. An SSB will be included in the 8.0 system as the following file: /usr/contrib/doc/SSB_8.0.

You should carefully read the latest version of *Installing and Updating HP-UX HP 9000* before starting an update to 8.0.

## Codewords for CD-ROM

Codewords are used with the update program to "unlock" software products stored on the CD-ROM disk in order to install or update these products on a specific computer system. Codewords are shipped to you from HP and are computed from information you supply with your order. This information includes the software product number that you have purchased and a unique hardware ID from the system you intend to install the software product on.

## Device File Names

You must know the device file name for your installation device (CD-ROM or DDS/DAT drive, cartridge or reel tape drive) so that you know whether or not you need to override the default that update uses.

- On a Series 300/400/700, update uses /dev/update.src as its default.
- On a Series 600/800, update uses /dev/rmt/0m as its default.

If this is not the device file for the installation device you intend to use, you must change the source device name once you are in the update program.

## Set TERM environment variable correctly.

If your TERM variable is not set properly, your display could behave strangely during an interactive update process.

## UH-3. SLIDE: Codewords and Hardware IDs

160

### Codewords and Hardware IDs

CODEWORDS:
- Enable address to licensed software
- Computed based on hardware ID, software products purchased, and CD-ROM disk part number
- Supplied by HP in one of two forms:
  - Codeword Certificates for pre-computed codewords
  - Codeword Entitlement Certificates for post-computed codewords

HARDWARE IDs:
- Must be a unique identifying number related to your system hardware
- Are electronically readable by software
- Examples:
  - HP-HIL ID module
  - SW-ID number (Series 700)
  - Station (link level) address for LAN interface (Series 300/400)
  - HP-IB CD-ROM serial number

51436 UH-3. 3 

## Student Notes

All HP software distributed on CD-ROM comes loaded on the disks shipped to you as part of your HP-UX Software on CD Media subscription service. This service will provide two CD disks, one containing the operating system and subsystems and the other containing HP applications. Most of these products can only be accessed using a special codeword that unlocks the products. The codewords for these disks may be:

- Pre-computed codewords
  Pre-computed codewords are provided on Codeword Certificates shipped with the CD disks.
- Post-computed codewords
  Directions for obtaining post-computed codewords are provided on Codeword Entitlement Certificates.

Whether or not codewords are pre-computed depends on how the order was placed. One factor is the hardware ID. If the order is placed specifying the hardware ID as some hardware other than the HIL ID module, the codewords cannot be pre-computed.

## UH-4. SLIDE: Preparing the File System for Updating

161

### Preparing the File System for Updating

- Make sure you have adequate disk space to contain the new release.
- Perform a full system backup.
- Clean up the HP-UX file system and check disks with `fsck`.
- Save a copy of your kernel to `/SYSBCKUP`
- Ensure that any files that affect system configuration are current.

51436 UH-4. 4 

## Student Notes

### Adequate Disk Space

If you are updating to HP-UX 8.0, you will *not* need more disk space then with the 7.0 release. Significant changes in HP-UX libraries have occurred that reduces the required space for HP-UX commands and HP supplied executables. These space savings were offset by increased functionality and new features, but not to the extent that more disk space is required. Of course if you are adding new applications, you must accommodate their space requirements.

Before you update your system, analyze the amount of space available on your destination disk. Make sure there is enough space to accommodate the update. Consider these factors:

- The files on your system.
- The size of your disk (or disks).
- The partitions and filesets you will select to load.
- The minimum free space required on your system's mounted file systems.

- Swap space requirements on your system.
- Future expectations for space requirements on your system.

The `update` program will tell you if you have insufficient disk space to update your system. However, if you suspect you might not have enough disk space to update your system, it is best to handle the problem before you begin. You can free disk space by removing files from your system or by creating symbolic links. You could also choose to mount another file system.

### Perform a full system backup.

It is essential to perform a full system backup before you update your system. This will allow you to restore any or all of your files if something goes wrong with the update.

### Clean up the HP-UX file system

You should start with a file system that is in good shape. You should have been periodically running `fsck` to check disks. It would be a good idea to do this before you start an `update` to be highly confident of the file system integrity.

### Save a copy of your kernel to /SYSBCKUP

If the update process recreates the kernel, this step is not necessary, since it will be done automatically for you. If however, the `update` fails for some reason, it's a good idea to be able to reboot your system from a "known good" kernel.

In most cases this will not be a "bad" thing to do. There are some problems with simply assuming you can reboot from the saved kernel in /SYSBCKUP. For example, when updating to another major release of the HP-UX the backup kernel may be an earlier revision and therefore incompatible with the rest of the HP-UX files and commands.

### Files that affect system configuration should be current

Ensure that any files that affect system configuration are properly backed up and correctly modified according to the installation instructions in your *README FIRST* documents. As we cover more about such files later in the course we will discuss various files that affect the system configuration, such as `/etc/checklist` and others.

## UH-5. SLIDE: Steps to Update HP-UX with Local Media

162

### Steps to Update HP-UX with Local Media

1. Execute `shutdown` to bring the system to a single-user state
2. Write protect your *UPDATE* media
3. Power on the update device, insert media, and wait for drive to ready
4. Copy update tools from media to disk
5. Execute `/etc/update` and follow the interactive menus

51436 UH-5. 5 

### Student Notes

This slide describes the steps to update you system software if you plan to use physical media such as tape or CD-ROM. If you plan to use a network server to update your system software across a network, then some steps will be different. The network update is covered separately in this module.

#### Execute `shutdown` to bring the system to a single-user state

Depending on the filesets that are being loaded, the system may or not be rebooted during the update process. Filesets that require kernel reconfiguration will reboot the system. If you unsure if you are updating this type of fileset, it is best to be prudent and bring your system to a single-user run-level with the `/etc/shutdown(1m)` command:

```
# cd /
# shutdown
```

We will talk more about shutting down the system later in the course. Do not perform this step if you intend to update across the network, since the command will terminate processes needed for LAN access.

### Write protect your *UPDATE* media

If you are not using CD-ROM to update, the media you are using is vulnerable if not write protected. For cartridge tape media, turn the write-protect arrow towards the "Safe" label on the tape. For the DDS/DAT tape media, push the write-protect tag towards the corner of the tape. For mag tape media, remove the write-protect ring on the back of the tape reel, if present.

### Power on the update device, insert media, and wait for drive to ready

Turn the power on for the update device. Then, after the device self test and warm-up are complete, insert the first *UPDATE* medium into the device. Now wait for the drive to become ready. Some devices will indicate readiness by turning on a light and others will turn off a light. Consult the hardware installation manual for your particular drive. Remember to put the device on "online".

### Copy update tools from media to disk

The update tools, the `update` program and its supporting files and programs, are written to support a particular kernel revision. Therefore, if updating for the previous version of the operating system to the current, the current update tools must be extracted from the media and downloaded to disk.

### Execute `/etc/update` and follow the interactive menus

Make sure you are in the root (/) directory before invoking `update`.

You should type `/etc/update` on the HP-UX command line to invoke interactive update process. The menus will lead you through the process of updating the system.



23c

## UH-6. SLIDE: Copying the Update Tools

163

### Copying the Update Tools

To copy the `/etc/update` program and other tools:

for Cartridge tape media:

```
# cd /
# tcio -iZ /dev/dev-file-name | tar -xvf - TOOL
```

for DDS/DAT and mag tape media:

```
# cd /
# tar -xvf /dev/dev-file-name TOOL
```

for CD-ROM media:

```
# cd /
# /etc/mount  /dev/bsrc  /UPDATE_CDROM  -r  -t cdfs
# tar -xvf /UPDATE_CDROM/TOOL
```

51436 UH-6. 6 

## Student Notes

There are different commands to copy the `update` program and other required tools from the update media. The command you use depends upon the type of media you are planning to use. In the commands given on the slide, you will need to supply the correct device file name for your update device. The device file name you use depends upon your particular system configuration. The default device files are usually present on your system, having been put there by the install process. The device file name you might use is:

- /dev/update.src for cartridge tapes.
- /dev/update.src for DAT tape.
- /dev/rmt/0m for reel tape.

Notice that the CD-ROM command is a different format than for other media types. Usually the location where the CD-ROM is mounted is /UPDATE_CDROM, but the location may vary on your system. If the CD-ROM drive is not mounted, it must first be mounted before you type this command.

After you execute the `tar` command to copy the `update` program on to your disk, you will have to wait several minutes for the first `TOOL` file to be extracted. Once the process begins, the files are extracted rapidly. The `TOOL` files are located at the beginning of the tape; however, `tar` is unaware of their location and as such will continue to search for these files until the end of the tape. To shorten this process, when the last `TOOL` file is echoed to display, wait 5-7 minutes to be sure it actually is the last `TOOL` file, and then press Break to end the `tar` process.



235

## UH-7. SLIDE: Update Main Menu

164

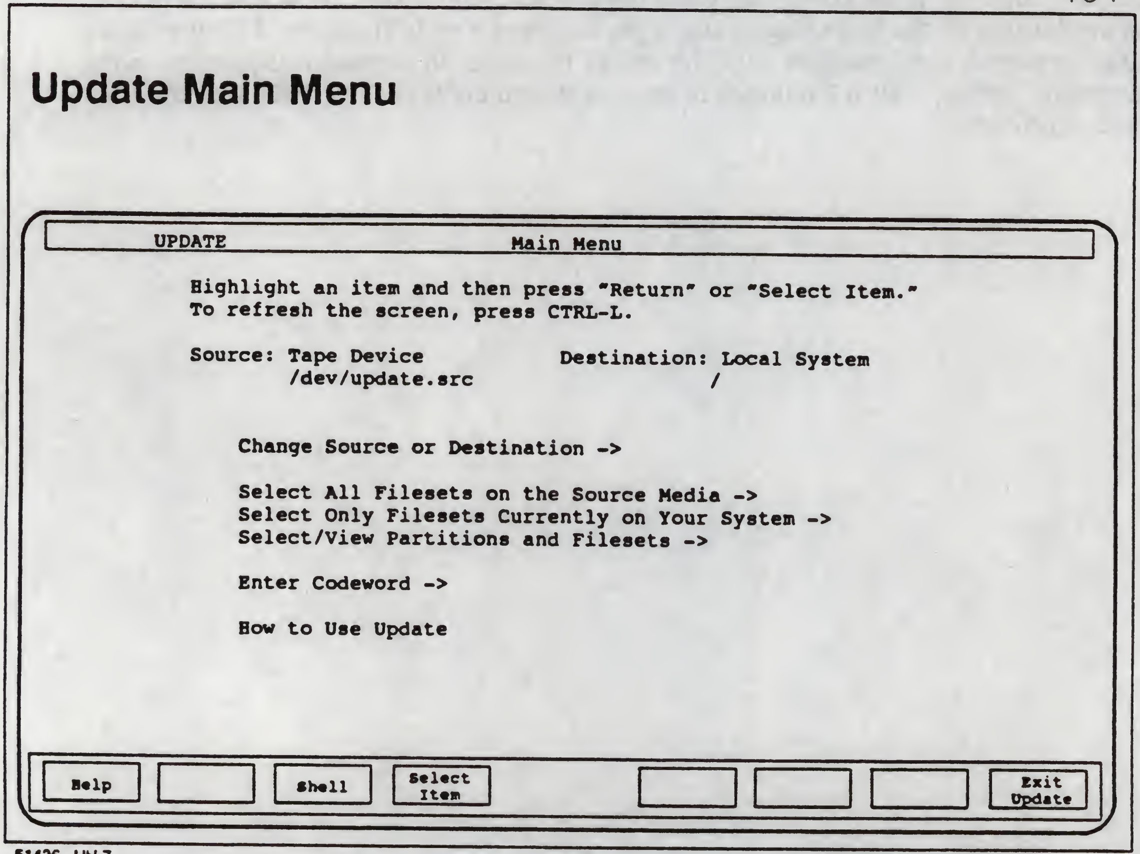


51436 UH-7. 7 

## Student Notes

### Navigating the Screen

Once you invoke `/etc/update`, you will see a message:

```
Initializing ...
```

followed by the screen as seen on the slide: Update Main Menu. You navigate in the `update` program the same way you navigate in SAM. To select an item:

- Use (Tab) or the cursor (arrow) keys to highlight the option you want to select.
- Press (Return) or `Select Item` to choose the current selection.
- Press (Tab) or (Shift)(Tab) to move forwards and backwards between the fields on the screen if a form is displayed..

If you are unsure about the function of certain menu items or data entry fields, highlight the item and then press `Help` for on-line information about the item. This capability applies only to those "update" screens where a softkey labeled "Help" appears.

If you want to quit the `update` process before going any further, press `Exit Update`. This capability applies only to those "update" screens where a softkey labeled "Exit Update" appears.

## Changing the Source or Destination

Select "Change Source or Destination" if your source or destination differs from the default listed on the Main Menu. A new menu will appear that allows you to enter the correct source and/or destination device. A sample of that menu (which is a pop-up type) follows in the next slide.

You would change the destination only if directed by the *README FIRST* notices shipped with your update media.

You may see several variations of the source specification:

```
Source: Tape Device             Destination: Local System
        /dev/update.src                      /

Source: Tape Device             Destination: Local System
        /dev/rmt/0m                          /

Source: Tape Device             Destination: Local System
        /UPDATE_CDROM                        /
```

## Which Filesets and Partitions To Select?

You should choose to "Select All Filesets From the Source Media." This ensures that you get all the files necessary for a proper update. If you decide you do not need some of the filesets after your update, you can delete a file set with the `rmfn` (remove function) command. This new command supplants the `sysrm` command.

If you choose to "Select/View Partitions and Filesets," you will get a menu listing all the possible partitions that you can load. You can use it to see which filesets belong with which partition. This menu is shown in a later slide.

Be aware that the contents of filesets and partitions have been rearranged from HP-UX 7.0. The names of partitions and filesets have also changed for the 8.0 release. Therefore, if you choose to "Select Only Filesets Currently on your System," there is a possibility that all files might not be updated.



232.

## UH-8. SLIDE: Changing the Source or Destination

t65

### Changing the Source or Destination

```
                              UPDATE Main Menu

        Highlight an item and then press "Return" or "Select Item."
        To refresh the screen, press CTRL-L.

        Source: Tape Device           Destination: Local System
              /dev/rmt/0m                           /

           Change Source or Destination ->

           +---------- Change Source or Destination -------------+
           |                                                     |
           |    Highlight an item and then press "Return"        |
           |    or "Select Item".                                |
           |                                                     |
           |    Update from source to destination:               |
           |                                                     |
           |    From Tape Device to Local System  . . .          |
           |    From CD-ROM (directory) to Local System . . .    |
           |    From Netdist Server to Local System . . .        |
           |                                                     |
           +-----------------------------------------------------+

 Help                        Select                                   Exit
                             Item                                     Window
```

51436 UH-8. 8 

## Student Notes

If you select the update Main Menu choice: "Change Source or Destination" a **pop up** menu window will appear over the Main Menu window. This is shown in the slide above.

If you are updating from a CD-ROM drive, you need to change the source device by selecting `From CD-ROM (directory) to Local System`. After selecting the menu choice, the source will be shown on the Update Main Menu as:

```
Source: Tape Device          Destination: Local System
        /UPDATE_CDROM                     /
```

The CD-ROM needs to be mounted in order for the update to be successful. Since the update tools are extracted from the CD-ROM, it may very well be already mounted. If you did not mount the required CD-ROM in the installation drive to the directory /UPDATE_CDROM before starting the update, you will not be able to do so inside the `update` program from the shell. You must exit the `update` program to mount the correct CD-ROM file system and then restart `update`.

```
# /etc/mount /dev/bsrc /UPDATE_CDROM -r -t cdfs
```

If you are updating from a remote system using the LAN, you need to change the source by selecting `From Netdist Server to Local System` instead of the default. We will cover "netdist" updates in separate slide near the end of this module.

## UH-9. SLIDE: From CD-ROM to Local System

166

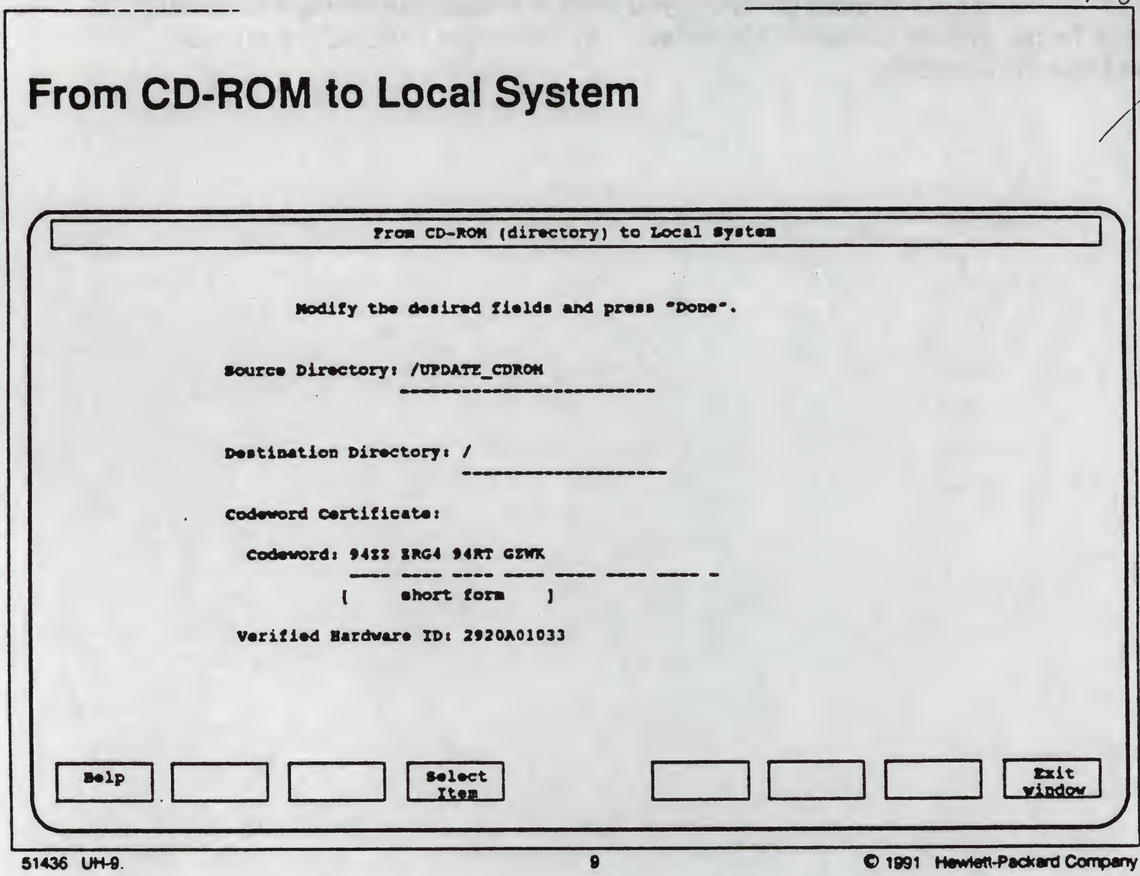


8.0

## Student Notes

When updating from codeword protected media such as CD-ROM disks, you must enter a codeword before selecting filesets. The codeword for your system software can be entered in the screen displayed here. The codeword you enter unlocks certain filesets on the "UPDATE" CD-ROM disk so that only the software products you purchased for this hardware can be installed on the hardware.

Tab down to the codeword field and enter either the 16 or the 29 character codeword you have been issued for this hardware. The codeword may be entered in either upper or lower case for letters. Do *not* enter spaces that may appear for legibility in the codeword on your software product certificate.

The code words are tied to hardware IDs. Possible choices for hardware IDs are:

- A machine ID number such as the SW_ID field in series 700 Stable Storage,
- A machine ID number from the LAN link level address on a series 300/400 system,
- A security number from an HP HIL 46084A ID module on a series 300/400/700.

- A security number from an HP C1707A HP-IB CD-ROM drive on a series 300/400.

You should not enter anything in the field "Verified Hardware ID." This field will be filled in by the `update` program if the codeword you supply matches one of the possible hardware IDs on your system. If the field is filled in and the window pops down before you can read it, just choose that option again. In the second go around, don't change anything, just select `Exit Window` to return to the "Main Menu" screen.

## UH-10. SLIDE: From Tape Device to Local System

167

### From Tape Device to Local System

```
                         From Tape Device to Local System

          Modify the desired fields and press "Done".

        Enter a "y" in the field to create or change the address of
        /dev/update.src or specify the name of a device to update from

        Create or change the address of       Address of current  source
        /dev/update.src? (y or n)   n              major number:     4
                                   ---             select code:      7
        Source:  /dev/update.src                   bus address:      0
                 ---------------                   unit number:      0
                                                 volume number:      0

    Destination Directory:
    ----------------------

 Help                      Done                                   Exit
                                                                  Window
```

51436 UH-10. 10 

## Student Notes

When updating from tape media, you are not required to enter a codeword.

Since we got here by choosing `From Tape Device to Local System`, you may interact with this screen to change the source or destination from the defaults displayed here. You do not have to enter anything. You complete the interaction with `Done` or `Exit Window`. `Exit Window` allows you to exit without making any changes even after you modified fields on the screen.

## UH-11. SLIDE: From Netdist Server to Local System 168

### From Netdist Server to Local System

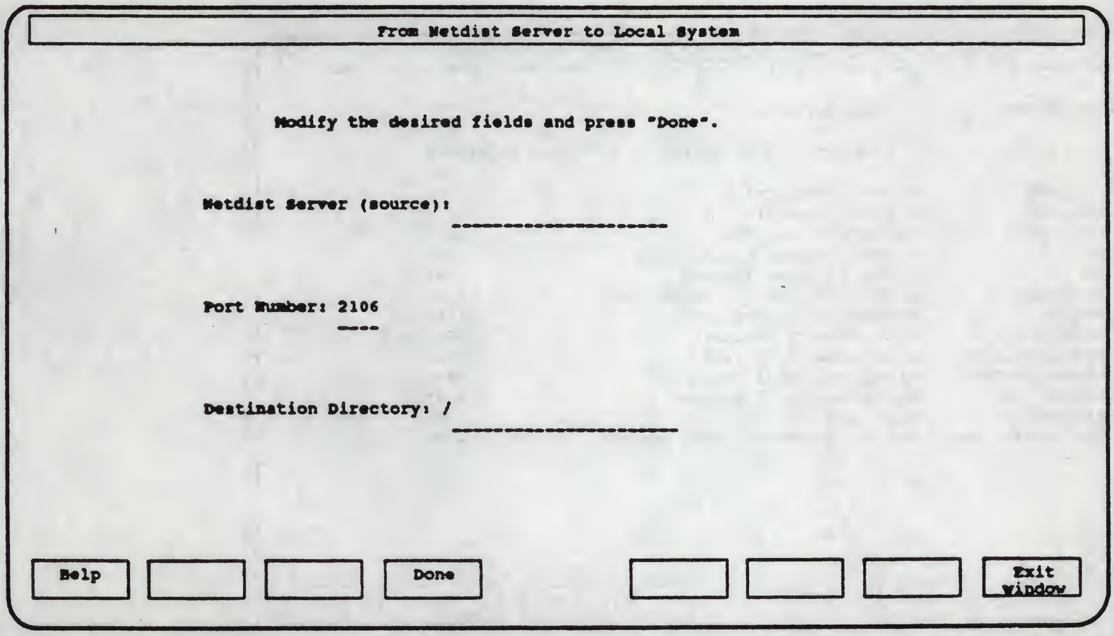


51436 UH-11. 11 

## Student Notes

### Netdist Server Selection

When updating from a Netdist Server, you may be required to enter a codeword.

Since we got here by choosing `From Netdist Server to Local System`, you must interact with this screen to identify the Netdist Server or to change destination from the default displayed on the screen. You must name a host that has been properly configured to act as a netdist server. (We will cover configuration of netdist servers later).

You complete the interaction with `Done` or `Exit Window`. The latter allows you to exit without making any changes even after you modified fields on the screen.



235

## UH-12. SLIDE: Update Individual Partitions Screen

169

### Update Individual Partitions Screen

UPDATE View or Select Individual Partition

Source: /UPDATE_CDROM          Destination: /

| Selected | Name | Partition Description | Size in Kbytes |
|---|---|---|---|
| n | DATABASE | Database Products | 22724 |
| n | GRAPHICS | Graphics Products | 12071 |
| n | NETWORKING | Networking Products | 9891 |
| n | NLIO | Native Language I/O Products | 26757 |
| n | NLS | Native Language Support | 592 |
| n | OS-ADMIN | Recommended Administration Cmds | 7980 |
| n | OS-CORE | Recommended System Core | 11650 |
| n | OS-FEATURES | Selectable OS Features | 18626 |
| n | PROG-LANGUAGES | Programming Languages | 10564 |
| n | REFERENCE-DOC | Reference Manual Pages | 5813 |
| n | SHARED-LIBS | Runtime Shared Libraries | 4322 |
| n | WINDOWS | Windowing Products | 22703 |
| v | vvvvvvvvvvvvvvvv | vvvvvvvvvvvvvvvvvvvvvvvvvvvvvvvv | vvvvvvvvvvvvvv |

Help | | Shell | Start Loading | Disk Space | View Filesets | Global Select | Main Menu

51436 UH-12.          12          

## Student Notes

If you choose to "Select/View Partitions and Filesets" from the Main Menu, then you will get this menu. A "working" message may appear at the bottom of the Main Menu screen for some time before this screen is displayed.

The time for the "working" message varies from a minute on some systems for a local update, but may range up to several minutes on remote (across the LAN) updates. You should be patient, a lot of information about the files you will be updating is being gathered from the update media or the netdist server.

Above the Source and Destination, you will see a message that says:

```
Mark "y" or "n".  To pick and choose individual filesets within a partition,
press "View Filesets".  A "p" means that some filesets have been selected
within a partition.  Press "Start Loading" when selection is complete.
```

If you mark "y" next to a partition, then you have chosen all the filesets within that partition. If you only want to choose some of the filesets in a partition, then highlight the partition and press "View Filesets." This will take you to a new menu that lists the filesets for that partition. Mark the individual filesets that you want with a "y" and then press `Done`. When you get back to Partition Menu, you will see a "p" (for partial) next to the partition.

The 8.0 (or later) version of `update` has much more knowledge of fileset dependencies than earlier versions. You may get warning messages if you select a fileset and then decide to "de-select" it. The warnings simply relate that you attempted to remove files that are needed by other selected filesets (There are hidden dependencies). This is actually a great help, but periodic warnings as you pick and choose filesets may be a bit distracting. The `update` program will remember your attempts to delete files and filesets it has warned you about. If you then remove the dependencies, the previous files will be automatically marked for removal. This is a bit of a surprise the first time it happens, but will prevent the need to circle back through the list multiple times.

In general you are much safer by loading everything (if you have disk space) and then using the `rmfn(1m)` command to delete software functions/products you do not want. Otherwise you will need to know which filesets to not load during the update.

Once you have marked all the partitions (and filesets) that you want to load, press `Start Loading`. Once you do this you will get several messages from the `update` program.

Next, you will get a "Calculating disk space requirements" message that flashes at the bottom of the screen. `update` will warn you if available disk space is insufficient. If this happens, a new screen is available that allows you to delete either some of the currently selected files you plan to update or existing disk files without leaving `update`. This will happen before you begin loading.

## Messages

In the case of a single update medium, such as CD-ROM or Netdist updates, you will get one last chance to confirm your intentions, or exit from the update process before loading begins. You will see a message such as:

```
The update should complete without additional attention.

Will load 52 filesets (5813 Kbytes), including 52 filesets (5813 Kbytes)
directly selected and 0 filesets (0 Kbytes) selected due to dependencies.
Review the log file, /tmp/update.log, afterward by:

1) typing "more /tmp/update.log";

2) finding the correct date and time heading;

3) looking for any messages that begin with ERROR or WARNING.


This is your last chance to change your mind about your selections.
Continue?  (y or n)
```

In the case of multiple media (tapes) you might get a message similar to the following:



236

```
 NOTE:  Filesets are selected that reside on multiple media.  You must
        change media at some point during the update.
In spite of these warnings or notes, you might want to continue.


Will load 79 filesets (25813 Kbytes), including 52 filesets (21403 Kbytes)
directly selected and 27 filesets (4410 Kbytes) selected due to
dependencies.


Review the log file, /tmp/update.log, afterward by:

1) typing "more /tmp/update.log";

2) finding the correct date and time heading;

3) looking for any messages that begin with ERROR or WARNING.


This is your last chance to change your mind about your selections.
Continue?  (y or n)
```

This means that you will need to change tapes during the update process. Once you choose to continue, `update` will start loading filesets from the first update tape to your disk. The next slide shows the screen you will get while `update` is loading files.

**Note**

The `update` program might reboot the system as part of the update process. The program reboots the system if you've selected a fileset (or filesets) that are flagged to indicate a reboot is necessary. If a reboot is necessary, you will be warned before loading begins. You have the option of proceeding or of exiting the `update` program.

## UH-13. SLIDE: Loading Partitions and Filesets (Multiple Media)

170

### Loading Partitions and Filesets (Multiple Media)

```
UPDATE           Loading Partitions and Filesets

  Loading fileset:  ALLBASE-MAN       From /dev/update.src

                          From medium number:  1 of 3

  This Medium:  Kbytes loaded:          8106  of   51059
(4%)
                Loading fileset:           3  of      36

  All Media:    Kbytes loaded:          8106  of  117072
                Loading fileset:           3  of      97

          Estimated remaining hours:minutes:          1:43

       Summary of Messages (also logged to /tmp/update.log)

            -- searching for next fileset --

 Help   Main   Shell  Perform          Disk   File Sys  Exit
        Menu          Task             Info   Info      Task
```

51436 UH-13. 13 

## Student Notes

This is the screen you will get while **update** is loading your files from the update media to the disk. It summarizes the activity of the **update** program by telling you which medium you are loading, how many kilobytes have been loaded, and which fileset it is on. When the program goes from one fileset to the next, you will see the message:

**-- searching for next fileset --**

This lets you know that it is looking for the next fileset on the update media.

Once all the filesets are loaded, your update is complete. If necessary, the system will reboot. Log in and look at the update logs.



237

## UH-14. SLIDE: Loading Partitions and Filesets (Single Media)

171

### Loading Partitions and Filesets (Single Media)

```
UPDATE                 Loading Partitions and Filesets  <--

          Loading fileset:  ALLBASE-MAN        from /UPDATE_CDROM

                   K bytes loaded:             106 of          5813
(1%)
                   Loading fileset:              3 of            52
                   Estimated remaining hours:minutes:          1:15


                   Summary of Messages (also logged to
/tmp/update.log)


              -- searching for next fileset --


 Help    Main     Shell   Perform            Disk    File Sys   Exit
         Menu             Task               Info    Info       Task
```

51436 UH-14. 14 

## Student Notes

This is the screen you will get while **update** is loading your files from the update medium to the disk. It summarizes the activity of the **update** program by telling you which medium you are loading, how many kilobytes have been loaded, and which fileset it is on. When the program goes from one fileset to the next, you will see the following message:

```
-- searching for next fileset --
```

Usually this message will flash at the bottom of the screen. The message lets you know that the **update** program is searching for the next fileset on the update medium.

Once all the filesets are loaded, your update is complete. If necessary, the system will reboot. Log in as normal and look at the update logs.

## UH-15. SLIDE: Summary of the Update Process

172

### Summary of the Update Process

- In Update Main Menu, select "Change Source or Destination"
- In "Change Source or Destination" Menu, select one of:
  - "From Tape Device to Local System"
  - "From CD-ROM (directory) to Local System"
  - "From Netdist Server to Local System"
- Interact with pop-up windows, give required information
- Return to the Update Main menu and select one of:
  - "Select All Filesets on the Source Media"
  - "Select Only Filesets Currently on your System"
  - "Select/View Partitions and Filesets"
- Mark the desired filesets and partitions to load
- Press **Start Loading**
- Wait for completion, change media as required.

51436 UH-15. 15 

## Student Notes



238

## UH-16. SLIDE: After the Update

173

### After the Update

- Check the `/tmp/update.log` file
- Check the `/etc/newconfig` directory
- Make a backup kernel (in some cases)
- Make a full backup
- Make a recovery system (Series 300/400/700 only)

51436 UH-16. 16 

## Student Notes

### The /tmp/update.log

The `/tmp/update.log` file contains a description of the events and any errors that occurred during the update process. The following items are message labels and their meanings. The actual messages are indented nine spaces.

The label `=======` indicates that a task within `update` is beginning or has completed. For example:

```
=======...BEGINNING UPDATE PROGRAM
```

The label `ERROR:` indicates that the program cannot proceed, or that it needs corrective action. In some cases this impacts `update` so much that it cannot continue. An example of a severe error is:

```
ERROR:  Destination directory "/mount" is invailid: No such file or directory.
```

The label `WARNING:` usually indicates the program can continue. However, something went wrong or requires attention, either now or later. One example that can often be ignored is:

```
WARNING: Cannot access /etc/checklist file: No such file or directory.
```

Another example of a Warning that is more serious and may prevent `update` from working as expected is:

```
WARNING: Cannot find any sources of hardware IDs for codeword verification.
         See the System Administrator Manual.
```

The `NOTE:` label indicates that something out of the ordinary or worth special attention has happened. The message may require no action on your part, in other cases the `NOTE:` message will require action. In some cases you must infer the action you must take after the update. For example:

```
        * Beginning to load fileset "TOOL".
NOTE:     Did not remove file "/etc/#update":  Text file busy (errno = 26).
        * Successfully loaded fileset "TOOL".
        * Beginning to load fileset "UX-CORE".
NOTE:     Did not remove file "/etc/#init":  Text file busy (errno = 26).
        * Successfully loaded fileset "UX-CORE".
```

You will have to remove the duplicate copy of the `update` and `init` programs which have been renamed: "/etc/#update" and "/etc/#init" after the update completes.

## Checking the /etc/newconfig Directory

The `/etc/newconfig` directory contains new versions of some system files normally put in the `/etc` directory (for example, `rc`, `brc`, `backup`). Because you might have edited the original versions of these files, they are not replaced by either `update` or update's customize scripts. Use the `diff` command to find the differences between your old files and new files placed in `/etc/newconfig`. Then you can incorporate your changes into the new versions of these files and move the new files into the `/etc` directory. Check the `README` file for information about the files in `/etc/newconfig`.

If you have problems with your update see the chapter, Updating HP-UX, of *Installing and Updating HP-UX* for your system type. Another good place to look is *Solving HP-UX Problems*, a new book at the 8.0 release.

## Make a Backup Copy of the Kernel

You only need to make a new backup kernel under certain conditions.

- If you are updating from one major revision of the OS to another, as in 7.0 to 8.0. In that case the kernel in /SYSBCKUP may not work with your new system files.
- The update did not make a backup kernel for you.

You can tell if you need to do this by reading the update.log file described above. The kernel will have been backed up during kernel rebuild when you see these messages in the log file:

```
          ...
NOTE:     Copied /etc/conf/dfile.min to /etc/conf/dfile.
```

```
NOTE:    Saved /hp-ux in /SYSBCKUP.
       * Generating a new kernel.
       * Following is output from "/etc/config":
         ...
```

If you see these messages you can be sure that the update process copied your kernel. Doing this again will overwrite the backup copy. You may want to overwrite the old kernel if you have updated from 7.0 to 8.0 for example. Usually this step is done prior to kernel reconfiguration which is covered in a separate module in this course.

### Make a Full Backup

This will be covered in a separate module in this course. The step is an important one, since you can always get your file system back to this point in the future should there be a serious problem with your disk or file system.

### Make a Recovery System

This will be covered later in this module. Making a recovery system applies only to Series 300/400/700 systems. The step is also an important one, since you can use the recovery system in the future to recover from a serious problem with your disk or file system.

## UH-17. SLIDE: Configuring for a Network Update

174

### Configuring for a Network Update

- Verify correct network operation between client and server
- Configure one host as an update server
  - Create `/netdist` directory
  - Using `updist`, load the software
  - Set up public network access
  - Put the `TOOL` fileset in a public directory
  - Start the network update server daemon, `netdistd`

51436 UH-17. 17 

## Student Notes

The server should have already been configured for correct LAN operation before configuring for as a network update server. Nevertheless, you should verify correct network operation on client and server. The specifics are beyond the scope of this class, but a cookbook version of the LAN setup process is described in an appendix to this course.

Configure one host as an update server. The server system needs a directory created to hold an image of the files from the distribution media. Sufficient disk space is required to hold the duplicate copies of all the files from the update media. This can be a substantial amount of disk space depending upon the software products expected to be supported for network update.

Steps to configure an update server:



240

### Create /netdist directory

The server system needs a root level directory `/netdist` created to hold an image of the files from the distribution media.

Ensure that sufficient disk space is available to hold copies of all the filesets from the update media plus additional space for overhead information about filesets, dependencies, sizes, etc. This can be a substantial amount of disk space depending upon the software products expected to be supported for network update. Check the README notices for sizes of the filesets. The filesets account for the lion's share of the disk space.

The `updist` program will help with an option to check disk space before actually uploading files.

### Using updist, load the software

The `/etc/updist` command will load the filesets from the **UPDATE** media onto the server's disk. The behavior of the `updist` program is similar to `update`, except that the files will be loaded into a special directory, and no customization is done after the process completes. At the completion of the 'updist'' program, the file system of the server remains unchanged except for the new (and very large) directory, `/netdist`.

The Main Menu of `updist` is very similar to the `update` main menu with a few minor changes:

All references to "Update" are now to "Updist"
Destination is now the "/netdist" directory vs the "/" directory.
All references to the destination will now be to a directory "/netdist" rather than the server's file system "/" as in `update`.

Follow the same steps to select desired filesets and partitions and the like as you would with `update`. Refer to the summary slide earlier for a review of the steps, mentally substituting `updist` for `update`.

### Set up public network access

A user login that allows read access to the filesets, but restricted access to the system is needed by client systems that wish to `update` using the netdist server. Because of the built-in security in the anonymous `ftp` user login, this account is a good choice for this type of access. The `ftp` login does not usually require a passwd be provided, hence the name "anonymous ftp". Setup of "anonymous ftp" can be done with a few key presses using SAM if your system is already up and running on the network.

The specific steps are recounted here. From the "System Administration Manager" main menu, select "Networks/Communications", then select "ARPA Services Configuration", and finally select "Create Public Account for File Transfers". Make sure you allow remote users to retrieve files when prompted by SAM. This will allow access to `/users/ftp/dist`.

### Put the TOOL fileset in a public directory

The `TOOL` fileset from the update media should be copied into `/users/ftp/dist` on the server. This allows the remote system to use the ARPA network service `ftp` to extract the update program from the server before running update.

For a server that intends to serve series 300 or 400 systems, in the directory: **/users/ftp/dist** you will put a copy of the TOOL fileset with the command:

```
# cd /netdist/300/TOOL/product
# tar cvf  /users/ftp/dist/TOOL.300  etc system
# chmod 444  /users/ftp/dist/TOOL.300
```

Replace 300 with 800 or 700 as appropriate. A network **updist** server can support all three platform's software if disk space is sufficient.

### Start the network update server daemon, netdistd

This cannot be done in the startup file, so you will have to type in the command manually:

```
# /etc/netdistd -l    # use a lowercase "L"
```

## UH-18. SLIDE: Performing a Network Update

175

### Performing a Network Update

- Verify connectivity to server system.
- Load the TOOL fileset using `ftp` to client.
- Run `update` on client system
- Complete update process as before

51436 UH-18. 18 

## Student Notes

Updating across the LAN is really just as simple as using local physical media to perform an update. You can only `update` your system using this method, you can not `install` a system across the network.

The steps required to perform a network based update.

### Verify connectivity to server system.

Log in to the remote system that is the designated netdist server to verify the connectivity, and ensure that the netdist daemon is running.

### Load the TOOL fileset using `ftp` to client.

Copy the Tools fileset across the network from the server. The steps from the client side are:

```
# ftp netdist-server
```



242

```
login:   ftp
passwd:  anonymous
ftp>  get dist/TOOL.300 /tmp/TOOL
```

... file is copied ...

*or* as appropriate to the type of the destination system

### Run /etc/update on client system.

Note that this is `update` on the client system, *not* `updist` which is used only in configuring the server system.

In Update Main Menu, select "Change Source or Destination"

In "Change Source or Destination" Menu, select: "From Netdist Server to Local System"

You will have to interact with pop-up windows, and give required information such as codewords, host names of the server, etc.

Return to the Update Main menu and select one of: "Select All Filesets on the Source Media" or "Select Only Filesets Currently on your System" or "Select/View Partitions and Filesets"

Mark the desired filesets and partitions to load

Press `Start Loading`

Wait for completion.

Complete update process as before Proceed to post-update procedures described above for the local media based update process.

# Module UH — Updating HP-UX

## UH-19. LAB: Updating HP-UX (Series 300/400)

### Directions

With your lab partner, perform the following tasks.

1. Use a current set of HP-UX *UPDATE* media to perform an update for a specific product. Check with your instructor to find out what filesets you should load.

2. Use a remote host configured as a netdist server to perform an update for a specific software product. Check with your instructor to find out what filesets you should load and the name of the remote server.

## UH-20. LAB: Updating HP-UX (Series 600/800)

### Directions

For this lab, simulations have been created for to allow you to gain experience with the update process.

1. Type `simulate` and select the update functionality.

# Module CA — Common Administration Tasks

## Objectives

Upon completion of this module, you will be able to:

- Explain the function of the System Administration Manager (SAM).
- Run SAM and navigate the menus.
- Use SAM to add users.
- Add users using the manual method.
- Maintain the `/etc/passwd` and `/etc/group` files.
- Identify required entries in `/etc/passwd` and `/etc/group`.
- Employ basic security procedures.
- Communicate with users on the system.
- Use the `at` and `batch` commands.
- Schedule programs for repetitive invocation with `cron`.

# Module CA — Common Administration Tasks

## CA-1. SLIDE: The System Administration Manager Roadmap

166

### The System Administration Manager

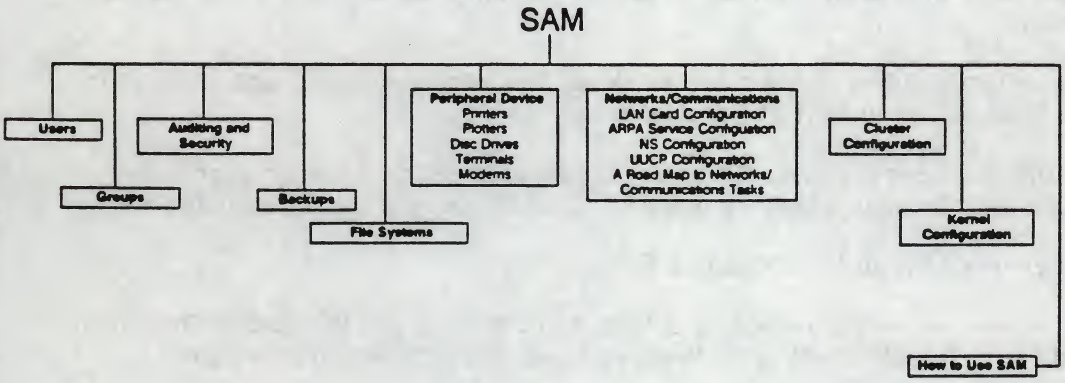


51436/51482 CA-1. 108 

## Student Notes

The System Administration Manager (SAM) is a menu-driven tool designed to perform typical system administration tasks. The slide shows some of the tasks SAM can perform.

Some of the benefits of using SAM are listed below:

- Instead of executing commands from a shell, you work through menus that guide task selection and facilitate data entry.
- Tasks are easier to perform because you need not remember (or type) complex commands.
- You get a rich set of functions, and those functions provide significant options and control.
- You can use SAM on Series 300 or Series 800 systems without relearning anything. The menus are the mostly the same.

While SAM may make it easier to perform certain tasks, the manual method of performing tasks is more flexible and more powerful. As a system administrator, it is very important that you understand the



245

manual way of doing things so that when you need to do something that SAM cannot do, you have the skills. The following items suggest a strategy for using SAM:

- Use SAM when you can to administer your system.
- Since SAM does not accommodate every task you need to perform, learn the manual procedure for performing a task.
- While performing a task with SAM, if you encounter a situation that SAM cannot accommodate, escape to a shell and perform the task manually.
- Use the manual procedure when SAM cannot perform a task or you know (as an expert) how you want to customize a functionality.

Remember, administering a system requires problem solving skills. The more you understand about your system, the better equipped you will be to solve the problems.

Listed below are the tasks that SAM can perform. (These are shown on the slide but are repeated for readability.) Users—-- Add a User Remove a User View/Modify User Deactivate User Reactivate User

Groups———Add a Group Remove a Group View/Modify a Group

Auditing and Security————————- Turn Auditing ON Set Audit Monitor Log Parameters Convert to Trusted System View Audit Logs View/Modify What is Being Audited: Users Events System Calls

File Systems———————Change File System Parameters Add a Remote (NFS) File System Remove a Remote (NFS) File System Convert File System to Long File Names Add a Hard Disk Drive Remove a Hard Disk Drive Change a Hard Disk Drive Address Add Dynamic Swap View File System Information View Disk Space Information

Peripheral Devices——————Printers and Plotters: Add a Local Printer Add a Remote Printer Enable a Printer Disable a Printer Remove a Printer Set the System Default Printer Shut Down the Spool System View Printer Status Information View Print Requests Add a Hard Disk Drive Remove a Hard Disk Drive Change a Hard Disk Drive Address View Disk Space Information Add a Terminal or Modem

Networks/Communications————————- LAN Card Configuration: Configure a New LAN Card View/Modify a LAN Card's Configuration Power Up a LAN Card Power Down a LAN Card ARPA Service Configuration: Add Connectivity to Remote System View/Remove Connectivity to a Remote System Specify the Default Gateway View/Modify Local Services' Security Let Remote Users Bypass Password Security Let Remote Users Become Super-User without a Password Enable Network Mail Create Public Account for File Transfers NS Configuration: Assign/Modify Your System's Nodename Add Connectivity to a Remote System Remove Connectivity to a Remote System Allow or Deny Access to Local Services UUCP Configuration: Add a Device Remove a Device View/Modify Device Configurations Add a System Remove a System View/Modify System Configurations Modify System Permissions A Road Map to Networks/Communications Tasks:

Cluster Configuration———————- Add Cluster Clients Remove Cluster Clients Create an HP-UX Cluster

Kernel Configuration——————Change I/O Configuration Modify Operating System Parameters Modify Primary Swap Partition Add/Remove Subsystems (NFS, LAN, NS, CD-ROM, etc.) Generate a New Kernel

How to Use SAM————————

## CA-2. SLIDE: How to Use SAM

167

### How to Use SAM

```
SAM                     System Administration Manager

        Highlight an item and then press "Return" or "Select Item."

        Users ->
        Groups ->
        Auditing and Security (Trusted Systems) ->

        File Systems ->
        Peripheral Devices ->
        Backup and Recovery ->

        Networks/Communications ->
        Cluster Configuration ->
        Kernel Configuration ->
        Task Customization. . .
        Other Utilities ->


        How to Use SAM ->

[Help] [    ] [Shell] [Select Item]    [    ] [    ] [    ] [Exit SAM]
```

51436/51482 CA-2. 109 

## Student Notes

To run SAM, simply log in as root and type:

```
# sam
```

The main menu is shown on the slide. An arrow at the end of an option means you get another menu. So, each of these selections takes you to a sub-menu. To select an item:

1. Use (Tab) or the arrow keys to highlight the option you want.
2. Press (Return) or `Select Item`.

Other function keys include:

| | |
|---|---|
| Perform Task | Executes a task that you have selected. |
| Refresh | Redraws your screen. |

| | |
|---|---|
| Help | Displays context-sensitive help, which contains additional information about using SAM. |
| Shell | Escapes to the shell. Type `exit` to leave the shell and get back into SAM. |
| Main Menu | Displays the main menu. |
| Exit Task | Exits the current menu or data entry screen. |

SAM has four types of screens:

- Menu screens show options (tasks) you can select.
- Data entry screens provide a form that has fields for entering data.
- Help screens show information about an item. Each highlighted item has a Help screen. Some Help screens list valid entries for a field.
- Feedback screens show messages from SAM:
  - Error messages appear when you type invalid data.
  - Progress messages indicate that SAM is performing a task.
  - Confirmation messages let you enter y or n to confirm taking an action or cancel it.
  - Termination status messages indicate a task has completed.



247

## CA-3. SLIDE: Adding Users with SAM

168

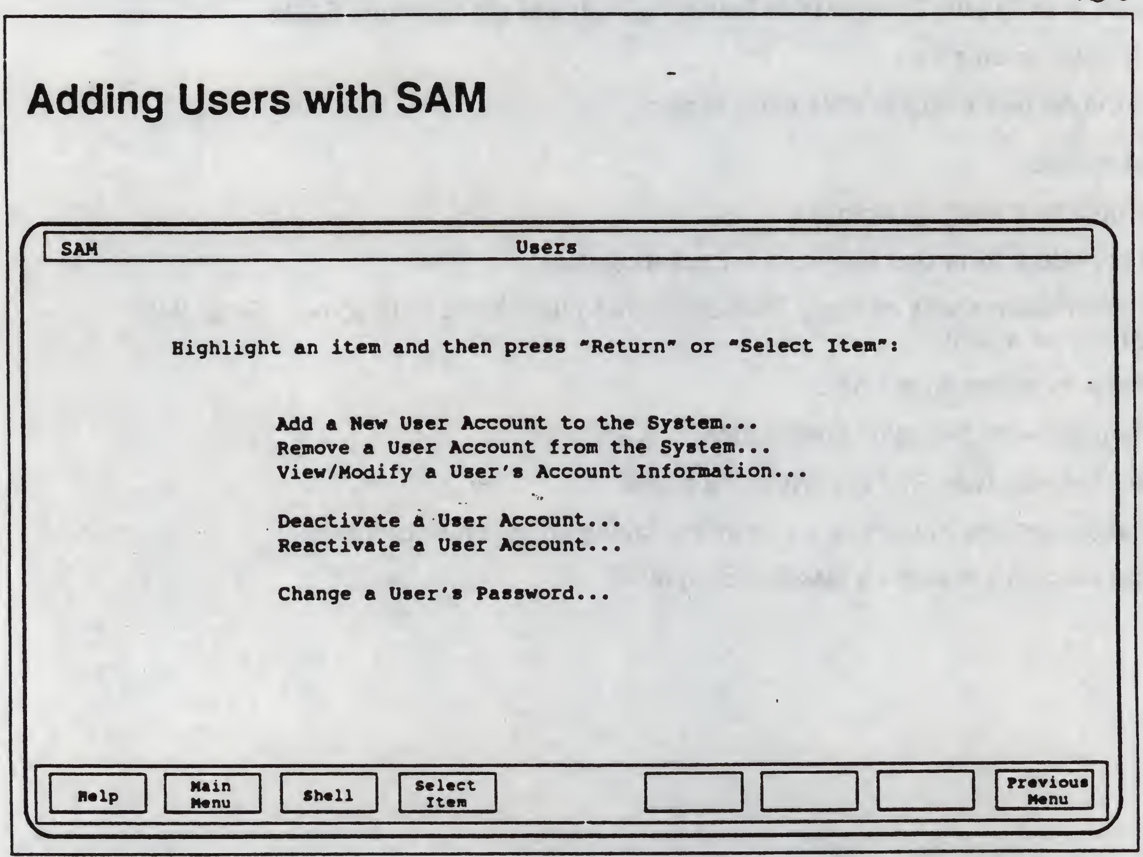


51436/51482 CA-3. 110 

## Student Notes

You will get this menu when you select *Users ->* from the Main SAM screen. *Add a New User Account to the System* should be highlighted. Press the (Select Item) key to select this item.

This screen also shows other user-related tasks that SAM can perform. If you deactivate a user, you deny that person access to the system until he/she is re-activated.

**Note** The ... after each menu item indicates that you will get another screen of information. The actual task will not be performed yet.

## CA-3. SLIDE: Adding Users with SAM (Continued)

169

### Adding Users with SAM (Continued)

```
SAM                  Add a New User Account to the System

        Fill in or modify the desired fields and then press "Perform Task":

               Login name......................________
               Primary group name..............users___
               Home directory..................../users/
               Start-up program................/bin/sh_
               Real name.......................__________(optional)
               Office location.................________ (optional)
               Office phone....................________ (optional)
               Home phone......................________ (optional)
               Modify Users Defaults? (y or n)  n

 Help | Main Menu | Shell | Perform Task |  |  |  | Exit Task
```

51436/51482 CA-3. 111 

## Student Notes

You will get this menu when you select *Add a New User Account to the System.* The highlighted areas are the places where you add information.

To add a user, just move the cursor around with (Tab), enter the new information, and press **Perform Task**.

Note that when you specify the user's home directory, the directories above the home directory must already exist, (for example, if you specify **/users/bugs** as a users home directory, **/users** must already exist on the system.

If you specify a user name that already exists, you will receive an error message.

We will talk more about these fields when we look at the **/etc/passwd** file.



248

If you answer yes to the "Modify user's defaults", you will get a second screen that will allow you to change this user's UID number.

## CA-4. SLIDE: Adding Users Manually

170

### Adding Users Manually

1. Add an entry to /etc/passwd
2. Add an entry to /etc/group
3. Create a home directory and set permissions to 755

```
# mkdir /users/bugs
# chmod 755 /users/bugs
```

4. Copy default setup files

```
# cp /etc/d.profile /users/bugs/.profile
```

5. Change the owner and group of the home directory and default files to that of the new user

```
# chown bugs /user/bugs
# chgrp users /users/bugs /users/bugs/.profile
```

51436/51482 CA-4. 112 

## Student Notes

The administrator of an HP-UX machine is responsible for maintaining the accounts of the users. This entails adding new user accounts when necessary, and removing or disabling the accounts of users who should no longer be permitted to access the machine.

To add a new user, there are only a few things to remember. Any new user account requires an entry in **/etc/passwd**. In addition, an entry for the new user must be made in **/etc/group**. We'll look at the format of the **/etc/passwd** and **/etc/group** files in more detail shortly.

A home directory must be created for the new user with the **mkdir** command. The permissions on this home directory should be 755 (drwxr-xr-x).

After creating the home directory, copy the default setup files to that directory. The particular setup files depend on the type of shell the user will be using. If the user will be using the Bourne or Korn shell:

```
# cp /etc/d.profile /users/bugs/.profile
```

If the user will be using the C shell:

```
# cp /etc/d.login /users/bugs/.login
# cp /etc/d.cshrc /users/bugs/.cshrc
```

The user may also have these other default files in his or her home directory. The system administrator can also set up these files to make it easier on the user:

- `.kshrc` (Korn shell startup file)
- `.exrc` (`vi` startup file)
- `.mailrc` (mailer startup file)

These startup files serve as a basis for allowing the new user to further customize their operating environment.

After copying the default setup files, the owner and group of the directory and setup files must be changed so that they are owned by the new user. (Since the home directory was created by root, root is the owner of the directory. Typically, the default setup files are also owned by root.) The **chown** and **chgrp** commands can be used to change owner and group information on directories and files.

**Note**

At 8.0 there is a new shell that uses all of the same files as the Korn Shell, it is called the **keysh**. This shell can also be specified when using SAM to set up a user's account.

Also at 8.0 the commands **chgrp, chmod, chown,** now have a -R option that allows you to use this command to recursively change a files owner or group down an entire tree structure.

## CA-5. SLIDE: The /etc/passwd File

171

### The /etc/passwd File

Syntax:

```
user_name:password:user_id:group_id:comment:login_directory:startup_program
```

Example:

```
root:t.jDsGls9ayKz:0:1:Super-user:/:/bin/ksh
bugs:,..:101:20:Bugs Bunny:/users/bugs:/bin/ksh
```

| | |
|---|---|
| *user_name* | Between one and eight characters |
| *password* | Encrypted, generated by system |
| *user_id* | Between 0-60000, 0 is root (0-99 are reserved for the system) |
| *group_id* | Associates a user with a group (any range) |
| *comment* | User information (such as full name and phone number) |
| *login_directory* | User's working directory after login |
| *startup_program* | Usually a shell, default is /bin/sh |

51436/51482 CA-5. 113 

## Student Notes

The slide shows what the fields in **/etc/passwd** mean. All fields are delimited by a colon (:). It is recommended that you make a backup copy of **/etc/passwd** before editing it, or that you use yank (yy) and put (p) within **vi**. More information on each of the fields is listed below.

- *user_name* is the login name which should be between one and eight characters. The first character should be alphabetic. If the name contains more than eight characters, only the first eight are significant.
- *password* is the encrypted password. It is generated by the system. The password should be six to eight characters, one of which should be numeric or a special character. If the password field is empty, then no password is associated with the login name. You should never leave the password field empty as it makes it very easy to break into a system.

  An asterisk (*) in the *password* field deactivates an account. Nothing you can type will encrypt to an asterisk, so, no one can login using the associated login name.



250

You can force a user to change their password right away the first time they login in by putting `,..` in the *password* field. (See Password Aging below.)

- *user_id* is the user ID (UID). UID zero (0) is reserved for root. Typically, values 1-100 are reserved for the system. Use numbers 101-60000 as UIDs for your users. Each UID must be unique. You can choose to use sequential UIDs for particular groups. For example:
  - 100-199 - marketing
  - 200-299 - engineering
  - 300-399 - managers
- *group_id* is the group ID (GID). This number corresponds with an entry in the `/etc/group` file.
- *comment* is the comment field. It allows you to add extra information about the users, such as the user's full name, telephone extension, organization, or building number. This field is used by the line printer spooler system and by the `finger` command.
- *login_directory* (or home directory) is the absolute path to the directory the user will be in when they log in. If this directory does not exist, or is invalid, then the root directory (/) is established as the user's current working directory.
- *startup_program* is the absolute path of a command to be executed when the user logs in. Typically, this is a shell. The shells that are usually used are `/bin/sh`, `/bin/ksh`, and `/bin/csh`. If the field is empty, the default is `/bin/sh`.

  The *startup_program* entry does not have to be a shell. For example, you could create the following entry in `/etc/passwd`:

  ```
  sync:rc70x.4,hGJ:20:1::/:/bin/sync
  ```

  The *startup_program* is `/bin/sync`. If you type `sync` at the `login:` prompt, and then type the appropriate password, the system will run the `/bin/sync` command and then log you out. If for some reason, you are forced to power off your system (for example, root cannot login), you could login in as `sync` to flush the buffers before powering down (and possibly prevent corrupting your disk). You could also create an entry in `/etc/passwd` that would shutdown your system. (UID should be 0.)

**Note** The permissions on the `passwd` file should be read only (`-r--r--r--`). When you edit the file with `vi`, use `:wq!` to save your changes.

## Using vipw

If you use `vi` to edit `/etc/passwd`, and a user attempts to change their password while you are editing, the user's change will not get entered into the file. To prevent this, you can use `vipw` when editing `/etc/passwd`.

```
# vipw /etc/passwd
```

This command puts a lock on the /etc/passwd file. If a user attempts to change their password, they will be told that the passwd file is busy.

After editing the /etc/passwd file, you can use the pwck command to find errors or inconsistencies in the file.

## Required Entries in /etc/passwd

```
root:jgR00Z280kf9g:0:3:root:/:/bin/ksh
daemon:*:1:5::/:/bin/sh
bin:*:2:2::/bin:/bin/sh
adm:*:4:4::/usr/adm:/bin/sh
uucp:*:5:3::/usr/spool/uucppublic:/usr/lib/uucp/uucico
lp:*:9:7:phantom for lp daeamon:/usr/spool/lp:/bin/sh
```

## Password Aging

You can control when a user must change his or her password with password aging. At the end of the encrypted password, you can add:

*,char1char2*

*char1* is the maximum number of weeks the password is valid and *char2* is the minimum number of weeks that must pass before the password can be changed. *char1* and *char2* may have the following values:

**Table CA-1.**

| Value | # of Weeks |
|---|---|
| . | 0 |
| / | 1 |
| 0-9 | 2-11 |
| A-Z | 12-37 |
| a-z | 38-63 |

As an example, if you wanted a user to change her password immediately, you would add the following to the end of her encrypted password in /etc/passwd:

```
,..
```

If you wanted the user to change his password sometime between 11 and 12 weeks, you would add the following to the end of his encrypted password:

```
,A9
```

Refer to *PASSWD(4)* for more information on the **/etc/passwd** file.

## CA-6. SLIDE: Changing Passwords

172

### Changing Passwords

- Users can change their own password (they must know their current password)

```
$ passwd
Changing password for bugs
Old password:
New password:
Re_enter new password:
$
```

- Root can change any user's password without knowing the current passwd

```
# passwd bugs
New password:
Re_enter new password:
#
```

- Root can change the root password

```
# passwd
Changing password for root
New password:
Re_enter new password:
#
```

51436/51482 CA-6. 114 

## Student Notes

Any time a password needs to be changed, whether by a regular user or the superuser, the **passwd** command is used. An ordinary user on the system can change his or her own password (but nobody else's). When invoked, the user will be prompted to enter the existing password. Upon entering the correct password, the user will be prompted to enter the new password. After the new password is entered, the user is prompted to enter it again. This is done for verification purposes and to ensure the user didn't inadvertently supply an extra character. If the second password does not match the first, the password is not changed and the user is returned to the shell.

If a user forgets his or her password, the user must seek the assistance of the administrator. The administrator can change any other user's password by invoking **passwd** with an argument of the user's login name.

If, for some reason the root password needs to be changed, the administrator should invoke **passwd** while logged in as the superuser.



252

If you are changing/assigning a password as a user:

- Passwords must contain at least 6 characters. Though a password may be assigned more characters, only the first eight are significant.
- Password must contain at least two alpha characters (upper or lower case) and at least one numeric or special character. This enforces a certain level of security within the password structure.

---

**Note**

When you use the **passwd** command, a copy of the old **/etc/passwd** file is saved in **/etc/opasswd**.

---

## CA-7. SLIDE: The /etc/group File

173

### The /etc/group File

*group_name:password:group_id:group_list*

Example:

```
other::1:root,daemon,uucp,who,date,games,sync
users::20:bugs,daffy,elmer,donald
```

| | |
|---|---|
| *group_name* | Group name |
| group_name | Group name |
| password | Typically unused |
| group_id | Number, corresponds to entry in /etc/passwd |
| group_list | List of users who may be members of the group |

51436/51482 CA-7. 115 

## Student Notes

The /etc/group file is used to define groups. The fields are delimited by a colon (:).

| | |
|---|---|
| group_name | is the mnemonic name associated with the group. If you do an `ll` on a file, you will see this name printed in the group field. |
| password | is typically not used, so it is blank. It can contain an encrypted group-level password if you implement privileged groups. |
| group_id | is the group ID (GID). This is the number that should be placed in the /etc/passwd file in the group_id field. This number is shared by all group members. |
| group_list | is a list of user names. These user names *are not* members of the group. Instead, this list defines the users who may become members of the group using the **newgrp** command. |

Look at the example below to see how the group_id is the same in the /etc/passwd and /etc/group files.

The /etc/passwd entry

```
bugs:,..:101:20:Bugs Bunny:/users/bugs:/bin/ksh
```

The /etc/group entry

```
users::20:bugs,daffy,elmer,donald
```

Use the grpck command to check the syntax of the /etc/group file after you edit it.

Note that a user can be a member of more than one group. A user can use the newgrp command to change to a different group.

```
$ newgrp group_name
```

## Required Entries in /etc/group

```
root::0:
other::1:
bin::2:
sys::3:
adm::4:
daemon::5:
mail::6:
lp::7:
```

For more information on the /etc/group file, see *GROUP(4)* in the *HP-UX Reference* manual.



254

## CA-8. SLIDE: Communicating with System Users

174

### Communicating with System Users

The news Command

- Displays files in /usr/news directory
- Used for long, less than critical messages

The mail, mailx, and elm Commands

- Mailer programs used to send messages to specific users
- Users read mail by invoking the specific mailer

The wall Command

- Sends a message to all users logged in
- Immediately interrupts whatever the user is doing

The /etc/motd File

- Contents are displayed by when a user logs in
- Used for short, important messages
- Typically changed on a daily basis

51436/51482 CA-8. 117 

## Student Notes

There are a number of ways that the system administrator can communicate with the users of an HP-UX system. The communication method used generally depends on the importance of the message that needs to be sent.

### The news Command

For messages that are not of great importance ("nice to know" as opposed to "need to know"), create a file containing the news item and place it in the /usr/news directory. When a user logs in, if there is a new entry in /usr/news, the user sees the following message:

```
news: news_filename
```

The user can then read the message with the **news** command.

Note that in order for the above to work, the following entry must be in **/etc/profile** for Bourne and Korn shell users:

```
if [ -f /usr/bin/news ]
then news -n
fi
```

For C shell users, the following entry must be in **/etc/csh.login**:

```
if ( -f /usr/bin/news ) then
   news -n
endif
```

## Mailers

If you need to send a long message to an individual user, then use one of the mailers.

## The wall Command

To simultaneously send a message to all users logged in, use the **wall** (write all) command. This command is typically used to generate a message that is of immediate concern to the users. For example, if, for some reason, the system must be shut down immediately, **wall** can be used to send a message to all the users currently logged in warning them of the impending shutdown. When invoked, **wall** reads standard output until an end-of-file is received and then sends the message as input to all logged in terminal lines.

```
# wall
The system will be shut down in 5 minutes.  Pleaselog off.
Ctrl-d

Broadcast Message from root (console) Sat Mar 18 11:22:43...
The system will be shut down in 5 minutes.  Please
log off.

#
```

**wall** can only be invoked by the superuser. When invoked, any permissions a user may have set to prevent someone from writing to their terminal (for example, **mesg n**) are overridden.

## The /etc/motd File

For messages that every user should be aware of, place an entry in `/etc/motd` (message of the day). For example, if you decide to shut the system down for an evening for preventive maintenance work, place a message in `/etc/motd` letting the users know which evening the system will be down so they can plan their work time accordingly.

The file `/etc/motd` is displayed every time a user logs in. The `/etc/profile` and `/etc/csh.login`, which are executed automatically during system startup, contain the command `cat /etc/motd`. This is covered in more detail in the System Startup module.

## CA-9. SLIDE: Security Considerations

175

### Security Considerations

- Set proper permissions on files and directories
- Protect passwords
- Log off unattended terminals
  - Set time out (TMOUT) variable if available
  - Use `lock` to lock an unattended terminal
- Monitor set-user ID programs
  - Use `ncheck -s` to produce a report of files with set-user ID mode on

uid gid st rwx rwx rwx

sticky bit.

51436/51482 CA-9. 118 

alleen 7.0

prog. loopt als owner.

## Student Notes

As an administrator of an HP-UX system, maintaining the overall security of the system is a primary responsibility. There are many aspects of system security, some of which we'll touch upon here; however, some preliminary words may be in order. The UNIX operating system, upon which HP-UX is based, was originally designed to be a research and development system which was no more security conscious than any other available operating system. Yet as UNIX and UNIX- derivative systems become more prevalent and their uses expanded beyond closed-shop research and development installations, there became a heightened interest in the security aspects of the operating system. Since the system administrator is responsible for the overall maintenance of the system, it became part of the administrator's job to employ and monitor security procedures.

prog. blijft altijd in swap area. => sneller.

als speciale perm. bit is 1 dan wordt de overeenkomstige x daar opgevraagd gezet, zie chmod.



256

### The User's Responsibility

The individual user should assume some responsibility for security. User's should:

- Set proper permissions on their files
- Protect their passwords
- Log out when they leave their terminal

Each individual user must be responsible for maintaining the correct permissions on their files and directories. Correct permissions are the only way to prevent someone from accessing unauthorized material. A prudent administrator may, on occasion, look at the permissions set on various important files and directories (for example, home directories) and advise their owners if something appears awry.

Each individual user must also be responsible for protecting his or her password. Users should be encouraged not to divulge their password to anyone, and should clearly understand the implications of supplying their password to another person.

Users should also be responsible for logging out of the system while their terminals are unattended. One way you, as an administrator, can enforce this is to use a time-out value such as the `TMOUT` (Bourn shell) or `autologout` (C shell) variable. With `TMOUT` or `autologout` set, if a login shell remains idle for a specified number of seconds, the shell is killed and the user is logged off.

Alternatively, a user can use the `lock` command to lock his or her terminal. When invoked, `lock` requests a key. The key must be entered twice for verification purposes.

```
$ lock
Key:
Again:
LOCKED
```

The terminal is locked until the key is entered again. This allows a users to leave temporarily without logging off.

### Tips for the Administrator

You should make sure that permissions on system files and directories are correct. The following list shows you what the permissions should be on some specific files and directories:

| | |
|---|---|
| `/etc/` | 755 (rwxr-xr-x), owned by `root`. |
| `/bin/*` | 555 (r-xr-xr-x), owned by `bin`. |
| `/etc/passwd` | 444 (r—r—r—), owned by `root`. |

Indiscriminate use of the set-user ID (or set-group ID) permission setting can result in some security difficulties. When a set-user ID program is executed, the effective user ID of the person executing the

program is changed to that of the owner of the program for the duration of the program's execution. For example, if the user john executes a program owned by bill which has the set-user ID bit set, then for the duration of the program's execution john's effective user ID is the same as bill's. Though this is an advantageous feature of the operating system, it does raise some concerns. For example, if a set-user ID program is writable, another program can be copied on top of it while retaining the set-user ID permission setting. Then, whenever the new program is executed, the effective user ID of the person executing the program is changed. Extreme caution must be taken with set-user ID programs that allow escaping to a shell. Since the effective user-ID of a user is changed when a set-user ID program is run, if a shell can be invoked from within the program, the spawned shell is established with the ownership being the same as the owner of the set-user ID program. There is a command that can be used to monitor set-user ID programs. The **ncheck -s** option will list all the files in a file system that have the set-user ID bit on.

rsh = restricted shell.



257

## CA-10. SLIDE: Other Security Considerations

176

### Other Security Considerations

- Keep current directory (.) out of root's `PATH` variable
- Insure your terminal protections deny access to others
- Monitor security logging files

| | |
|---|---|
| `/etc/wtmp` | logs successful login attempts |
| `/etc/btmp` | logs unsuccessful login attempts |
| `/usr/adm/sulog` | logs use of `su` command |
| `/etc/securetty` | specifies the `tty` files on which the root can log in (Series 800 only) |

51436/51482 CA-10. 119 

## Student Notes

If the root PATH variable contains an entry for the current directory, you have a potential security problem. The following example explains why.

Assume this is root's PATH

```
PATH=.:/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin:/etc
```

A user could create a script in his or her home directory (or any other directory) with the same name as a common command (for example, `$HOME/ls`). Assume the script contains the following commands

```
echo bandit::0:1:I am Superuser Now:/:/bin/ksh >> /etc/passwd
/bin/ls
```

If the user can convince the superuser to log in as root, change to his or her home directory, and run the `ls` command, then the `bandit` entry will get added to the `/etc/passwd` file. This happens because the

entry for current directory comes before /bin in root's PATH variable. (HP-UX searches for commands in directories in the order that they appear in the PATH variable.) So, the `ls` script in the user's home directory will execute instead of the real `ls` command in the /bin directory. Since root is executing this script, the `echo` command which appends a line to `/etc/passwd` will work. Then, since the *$HOME*/`ls` script invokes the real /bin/ls command, it will execute and everything will seem normal. Root wouldn't even know anything happened. Therefore, ***never*** put the current directory in root's PATH.

Second, insure that all terminal file protections are `crw--w----` and *not* `crw--w--w-`. If your terminal has write access by other terminals and your logged in as root, it is possible to program the softkeys from the remote terminal by writing a series of escape sequences to it. Then someone could also send escape sequences to execute the command programmed into the remote softkey with the capability of the user logged in at that remote terminal - that being root in this case. Its complicated, but the moral is: Keep the protections of your terminal assigned only to you if security is an issue. Also, while logged in as root, `mesg n` takes away write permissions to your terminal.

Third, monitor the security logging files. This will help tell you if anyone is trying to break system security. The files are below:

`/etc/wtmp` The system uses this binary file to keep a history of logins, logouts, and date changes. The system automatically creates this file, which grows without bound. Check the file regularly and empty it. Use `/etc/last` to access the contents of the file.

`/etc/btmp` If this binary file exists, the system uses it to keep track of *bad* login attempts. To get the information, you must create this file by entering:

```
# touch /etc/btmp
```

The file grows without bound, so check and empty it regularly. Use `/etc/lastb` to access the contents of the file.

**Note** The permissions on /etc/btmp should be 600.

`/usr/adm/sulog` If this text file exists, the system uses it to keep track of information on the use of the `su` command. It contains old and new user names, the corresponding terminal name, the date and time, and whether the attempt was successful. To enable logging of this information, create a zero-length file called `/usr/adm/sulog` with the `touch` command. To look at the information, use the `more` or `cat` command.

`/etc/securetty` If it exists, this text file specifies the `tty` files on which the root user can log in. You must explicitly create this file and place the `tty` device file names in it. To look at the information, use the `more` or `cat` command.

ACL ; ~~access~~ accounting / control list → per user protection

'secure system' ; C 2 (mate van beveiliging)

C 2
C 1
B 3
B 2
B 1
A 3
A 2
A 1

C 3 = standard unix.



25-8

**Table CA-2.**

| | enable logging* | display log file |
|---|---|---|
| /etc/wtmp | (automatic) | last |
| /etc/btmp | touch /etc/btmp | lastb |
| /usr/adm/sulog | touch /etc/adm/sulog | more /etc/adm/sulog |
| /etc/securetty | vi /etc/securetty | more /etc/securetty |

**Note** `/etc/wtmp` and `/etc/btmp` grow without bound. Empty them periodically. You can do this by removing the existing file with `rm` and creating a new zero-length file with `touch`.

## CA-11. SLIDE: Scheduling Programs for One Time Invocation

177

### Scheduling Programs for One Time Invocation

- at schedules a program to run at a specific time

```
$ at 3:00 pm Mar 3
echo "Birthday Time" > /dev/console
Ctrl-d
$
```

/usr/lib/cron/at.allow
" /at.deny

- batch executes a program in batch mode

```
$ batch
nroff file.frc > file.fmt
Ctrl-d
$
```

51436/51482 CA-11. 120 

## Student Notes

Frequently, it is desirable to execute a program at a later time, or to schedule a program to run at specific time. For example, if a backup program needs to execute everyday at 4 am, there should be mechanism to schedule the program to run so that an individual need not be there to directly invoke the program. HP-UX supplies three mechanisms to schedule programs:

- `at` - schedule a job to run "at" a specific time.
- `batch` - schedule a job to run in the background; the system decides when to run the job
- `cron`

`at` and `batch` are generally used to schedule a one- time invocation of a command at a later time, whereas `cron` is typically used to schedule commands to run at regular intervals.

Access to the `at` command is controlled by the files:

- `/usr/lib/cron/at.allow`

- `/usr/lib/cron/at.deny`

First the `at` command checks for the existence of the `at.allow` file. If `at.allow` exists, then permission to use `at` is granted to those users whose name is in the file. If `at.allow` does not exist, `at` checks for the existence of `at.deny`. If `at.deny` exists, then permission to use `at` is denied to those users whose name appears in the file. If both files exist, precedence is given to `at.allow` since `at` checks for the existence of this file first. If neither file exists, only root can use `at`.

If you create an empty `at.deny` file, you grant usage to all users. Similarly, if you create an empty `at.allow` file, you deny usage to all users.

The format for both files is one user-id per line.

## CA-12. SLIDE: Scheduling Programs for Repetitive Invocation

178

### Scheduling Programs for Repetitive Invocation

1. Ensure `cron` process is running

   `$ ps -ef | grep cron`

2. Root must add login name to `/usr/lib/cron/cron.allow`

   `# vi /usr/lib/cron/cron.allow`

3. User creates his/her `cronfile` containing programs to be submitted

   `$ vi cronfile` in home dir.

4. User submits `cronfile` to `cron` process with `crontab` command

   `$ crontab cronfile`

51436/51482 CA-12. 121 

## Student Notes

The `cron` process is normally started during boot-up by the `/etc/rc` script. Therefore, as long as you are in multi-user state, `cron` should be executing on your system. If for some reason it is not, start it running by typing:

```
# cron
```

Regular users, as well as root, can utilize `cron` for repetitive execution of programs. Jobs are submitted to `cron` with the `crontab` command. Root controls who can use `crontab` through the `/usr/lib/cron/cron.allow` file. Users are permitted to use the `crontab` command if their names appear in the `cron.allow` file. If `cron.allow` does not exist, then `/usr/lib/cron/cron.deny` is checked to determine if the user should be denied access. If both exist, `cron.allow` takes precedence. If neither file exists, only root is allowed to submit a job. An empty `cron.deny` file allows all to use `crontab`.



260

Once you have access to the crontab command, create a file called cronfile in your home directory. This file contains the commands you wish to submit to cron and the times you want them to be executed. We will look at the format of this file on the next slide.

To submit your job to cron, type:

```
$ crontab cronfile
```

This copies your local file called cronfile to the system cron directory /usr/spool/cron/crontabs/*user_name*. The file will be named after your user name. Once you submit your file with the crontab command, your job is activate.

Whenever cron executes a command, it writes a record to /usr/lib/cron/log (unless otherwise specified). This file grows without bound. It should be checked and emptied periodically.

### Files That Grow Without Bound

We have already seen several other files that grow with out bound, including /etc/wtmp and /etc/btmp. Files such as these should be monitored by the system administrator and appropriate action should be taken on a regular basis. You could do this with the cron facility. Below is a list (not necessarily complete) of files that tend to grow without bound.

vogt requests for am cron.

cron - f load requests zien

$\boxed{\text{Tmax}}$

**Table CA-3.**

| File Name | Written to by | Read by |
|---|---|---|
| /etc/wtmp | /etc/login | /etc/last |
| | /etc/init | |
| | /usr/lib/acct/accton | /usr/lib/acct/accton |
| /etc/btmp | /etc/login | /etc/lastb |
| /usr/adm/sulog | /bin/su | |
| /usr/lib/cron/log | /etc/cron | |
| /usr/spool/mqueue/syslog | /usr/lib/sendmail | /usr/bin/mailq |
| /usr/spool/mqueue/sendmail.st | | |
| /usr/spool/uucp/diallog | Many of the UUCP | |
| /usr/spool/uucp/errlog | commands such as | |
| /usr/spool/uucp/logfile | /usr/lib/uucp/uuxqt | |
| /usr/spool/uucp/syslog | /usr/lib/uucp/uucico | |
| /usr/spool/uucp/culog | | |
| /usr/spool/lp/log | /usr/lib/lpsched | /usr/bin/lpstat |
| | /usr/bin/lp | |
| /usr/adm/messages | /etc/dmesg | |

## CA-13. SLIDE: cronfile Format

(79

### cronfile Format

- Consists of lines of six fields each. The first five specify the following:
  - minute(s) after the hour (0-59)
  - hour(s) of the day (0-23)
  - day(s) of the month (1-31)
  - month(s) of the year (1-12)
  - day(s) of the week (0-6 with 0=Sunday)
- An "*" indicates all legal values
- Fields are separated by spaces or tabs
- The last field is the program to be executed
- Example:

```
0 * * * *          /bin/date  >  /dev/console
* 0-23 * * *       /bin/who  >>  /users/bugs/whofile
0 1 * * *          /etc/last  |  lp
0 6 * * 1,3,5      /etc/lastb |  lp
```

51436/51482 CA-13. 122 

## Student Notes

The crontab command allows you to submit jobs to cron to be executed at a later date and time. The commands and their associated execution times are entered in a cronfile.

The entries in cronfile must be in a specific format to be interpreted successfully by cron. Each entry in the file is a line containing six fields separated by white space or tabs. The first five fields contain integers which represent the date and time a command is to be executed. They are shown on the slide. Each of these fields may contain an asterisk which represents all legal values, or a list entries separated by commas. Each entry may be either a number or two numbers separated by a dash which specifies a range.

The last field is a string that is executed by the shell at the specified times. A percent character in this field (unless escaped by \) is translated to a new-line character.

You could use cron to help regulate your backups. An example is shown below:

```
0 23 * * 0      full backup command
0 4 * * 2-6     incremental backup command
```

**Note** You must redirect the standard output and standard error of your commands. If you do not do this, any output generated will be mailed to you.

In /etc/newconfig staat default crontab.

(cron wordt niet na install. gestart)

/etc/ dmesg - >> /usr/adm/messages.

toont laatst gequeede system ~~[illegible]~~ messages.



262

## CA-14. SLIDE: Modifying Your cronfile

180

### Modifying Your `cronfile`

1. List contents of existing `cronfile` (First move to the directory where your `cronfile` resides)

   `$ crontab -l > cronfile`

2. Make changes to `cronfile`

   `$ vi cronfile`

3. Re-submit your `cronfile` to `cron`

   `$ crontab cronfile`

51436/51482 CA-14. 123 

## Student Notes

Each user who is authorized to issue jobs to `cron` has *one* file in the directory `/usr/spool/cron/crontabs`. The name of this file is the user's login name. Once the user's `cronfile` is activated with the `crontab` command, any future use of this command will cause the crontab file in `/usr/spool/cron/crontabs` to be replaced.

To change information in your crontab file, retrieve the existing crontab file, modify it, and then re-submit the new crontab file. `crontab` will replace the old file with the new one.

Note that you use the `crontab -r` command to remove `cron` jobs.

## CA-15. WORKSESSION: Review Questions

1. Describe the seven fields of the `/etc/passwd` file.

2. Describe the fields of the `/etc/group` file.

3. What does it mean to set up groups? Explain.

4. Name the file steps to manually add a user.

5. Describe two ways to create a user and send him straight into an application program.

6. What is the difference between `/etc/profile` and `.profile`?

7. What is the `wall` command and how could it be useful?

8. If you were to shutdown the system tonight at midnight, how might you communicate this to all of your users?

9. What type of activities might be useful to schedule with cron?

# Module CA — Common Administration Tasks

## CA-16. LAB: Lab Exercises

### Directions

For some of the tasks below, you will have to take turns performing them with other members of your class. Only one person at a time should use SAM or modify the `/etc/passwd` file. Your instructor may choose to pass around a "tolken" to help make sure only one person at a time is performing certains tasks.

Or if p-shell is available you can run the commands from there.

1. What steps is SAM taking when performing the following tasks.

   Adding a new user

   De-activating a user

   Modifying a user's information

   Adding a new group

2. When it is your turn, invoke `sam` and add a user to your system. (You must be superuser to invoke SAM.) Use your name as a user name. Assign the user to a group called "class" and give him or her the Korn shell.

Now, exit SAM and look at the `/etc/password` and `/etc/group` files. Do you see the user you added? (DON'T FORGET TO PASS THE TOLKEN.)

3. What types of things can SAM *not do* with respect to adding users to the system ?

4. When it is your turn, add a user to the system using the manual method. This time, use your partner's name as the user name. (If you don't have a partner, pick any name.) Use a group called "class" and give the new user the C shell. Also, set the password so that the user must change the



264

password the first time he or she logs in. The steps are provided below. (DON'T FORGET TO PASS THE TOLKEN WHEN YOU ARE FINISHED EDITING THE passwd and group files.)

Add entries to /etc/passwd and /etc/group.

Create a home directory for the user and set the permissions and ownership correctly.

Copy default set up files to the home directory and set correct permissions and ownership.

Now, log in as the user you created. Did your password aging work?

5. Run the commands to check the integrity of the /etc/passwd and the /etc/group files. Discuss your findings with the instructor.

6. Log off of your session by typing in exit. At the login prompt, type in who. Now try date. What happens ? Why?

7. Try changing the password on your account to your first name. Does this work ? Why or why not?

8. Time schedule a command which runs at lunch time and echoes out to your terminal - "Lunch time - Where are we going ?"

9. Time schedule the `date` program to execute every minute with the display coming to your screen. (Hint: Make sure you redirect the standard output to your screen device.)

10. Now add another process to your time scheduled list. Have this one execute a program to display to your screen who is logged in. This program should run every 10 minutes.

# Module LP — Managing the LP Spooler

## Objectives

Upon completion of this module, you will be able to:

- Activate the LP (line printer) spooler.
- Add a serial printer to your system.
- Add a printer to the LP spooler system.
- Start and stop the LP spooler.
- Remove a printer or class of printers from your system.
- Check LP spooler status.
- Enable and disable a printer class.
- Set the default printer.
- Move printing requests to other destinations.
- Understand what print classes are.
- And understand what priority fences are.

# Module LP — Managing the LP Spooler

## LP-1. SLIDE: An Overview of the Spooling System

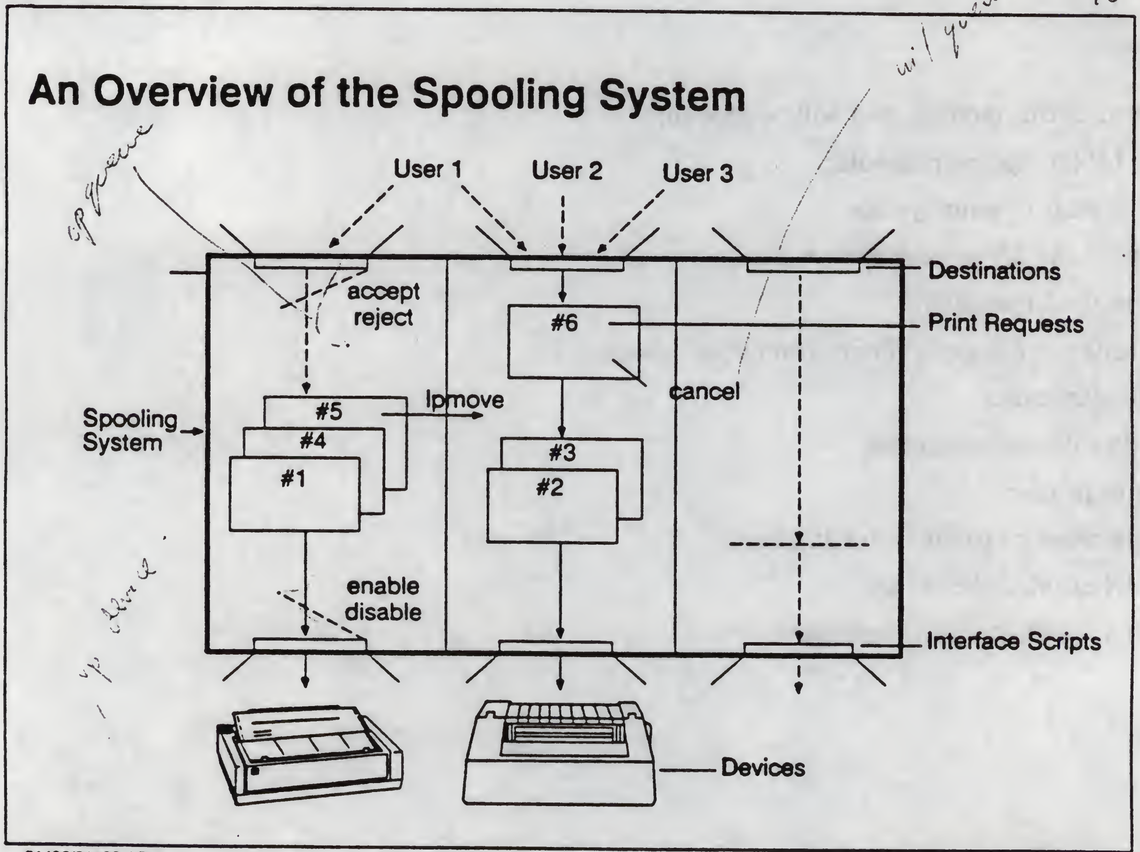


51436/51482 LP-1. 168 

## Student Notes

On a multi-user system, access to the printers requires careful management and control. Since users could possibly send print requests to the same printer at the same time, there needs to be a way to:

- Make sure each file is printed separately.
- Determine which file will be printed first.

The **LP spooler** is a collection of utilities and commands that controls the print requests of users. The LP spooler stores print requests in the spool directory until a printer is available. When a printer is available, print requests are processed, one at a time, in the background.

The spooling system can be customized by the administrator. For example, the administrator can choose where requests are spooled, or the administrator can also group several printers into a **class** of printers which can increase the overall efficiency of the system.

# Module LP — Managing the LP Spooler

After the LP spooler system is installed, any user can submit a job to be printed, obtain the status of all or any printer, cancel any print job, or declare printers to be in and out of service.

New features of the spooling system at 8.0 include:

- Support for centronics printers, the driver is `lpr2`.
- Support of laserjet IIP and laserjet III.
- Spooling is also now available on cnodes.
- The spooler can be invoked with remsh.

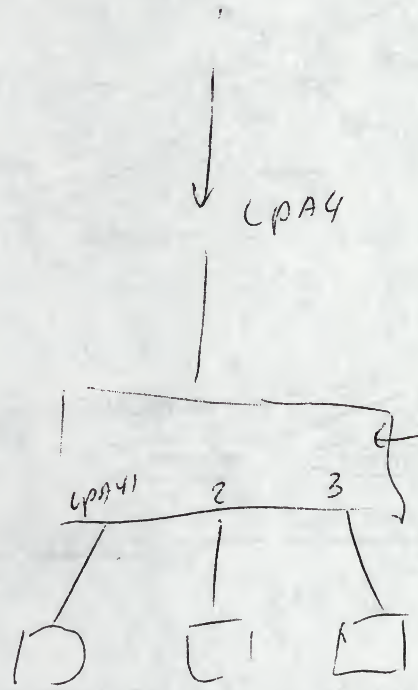


spooler zoekt uit welke prn vrij is

## LP-2. SLIDE: Directory Overview

182

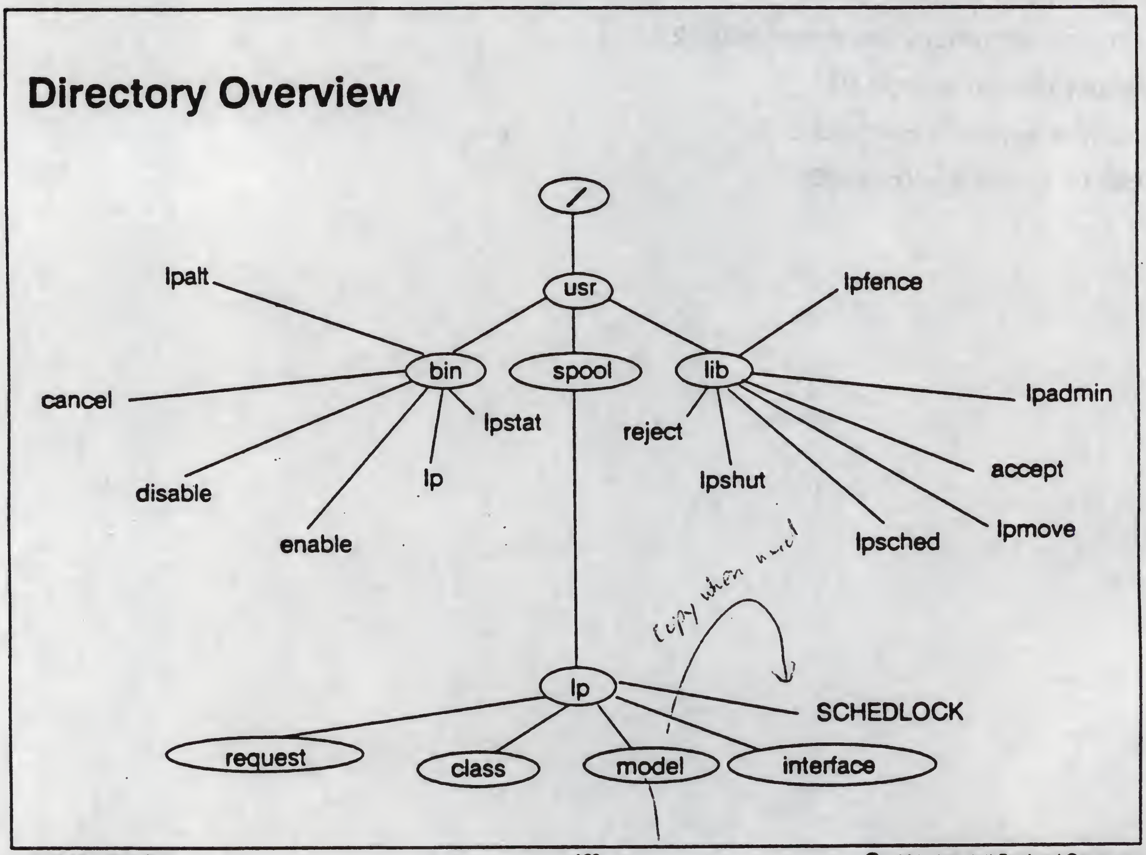


51436/51482 LP-2. 169 © 1991 Hewlett-Packard Company

printer specific scripts.

## Student Notes

The following is a road map of your way through the directories of the spooling system.

| | |
|---|---|
| `/usr/spool/lp` | LP spooler system parent directory. All information about the setup and printing queues is located here. |
| `/usr/spool/lp/class` | Printer classes directory. This contains the files that define how printers are grouped. |
| `/usr/spool/lp/model` | System-supplied interface programs. This directory contains the model shell scripts designed for particular printer models. |
| `/usr/spool/lp/interface` | Interface programs in use on your system. This has shell scripts from `/usr/spool/lp/model` that may be modified for particular printers. If |

| | |
|---|---|
| | interfacing a printer for which there is no model file, you may need to create an interface program for it. |
| `/usr/spool/lp/request` | Destination queues. This is where all `lp` requests are queued. It usually contains a subdirectory for each printer configured on the system. |
| `/usr/bin` | Contains user-executable commands, such as LP spooler commands that general users can execute. |
| `/usr/lib` | Contains administrator commands, such as the LP commands that only root or `lp` can execute. |
| `/usr/spool/lp/member` | Lists all printers, one file per printer. |
| `/usr/spool/lp/fonts` | Contains fonts for Laserjet printers. |



268

## LP-3. What Happens When a File is Submitted With `lp`

133

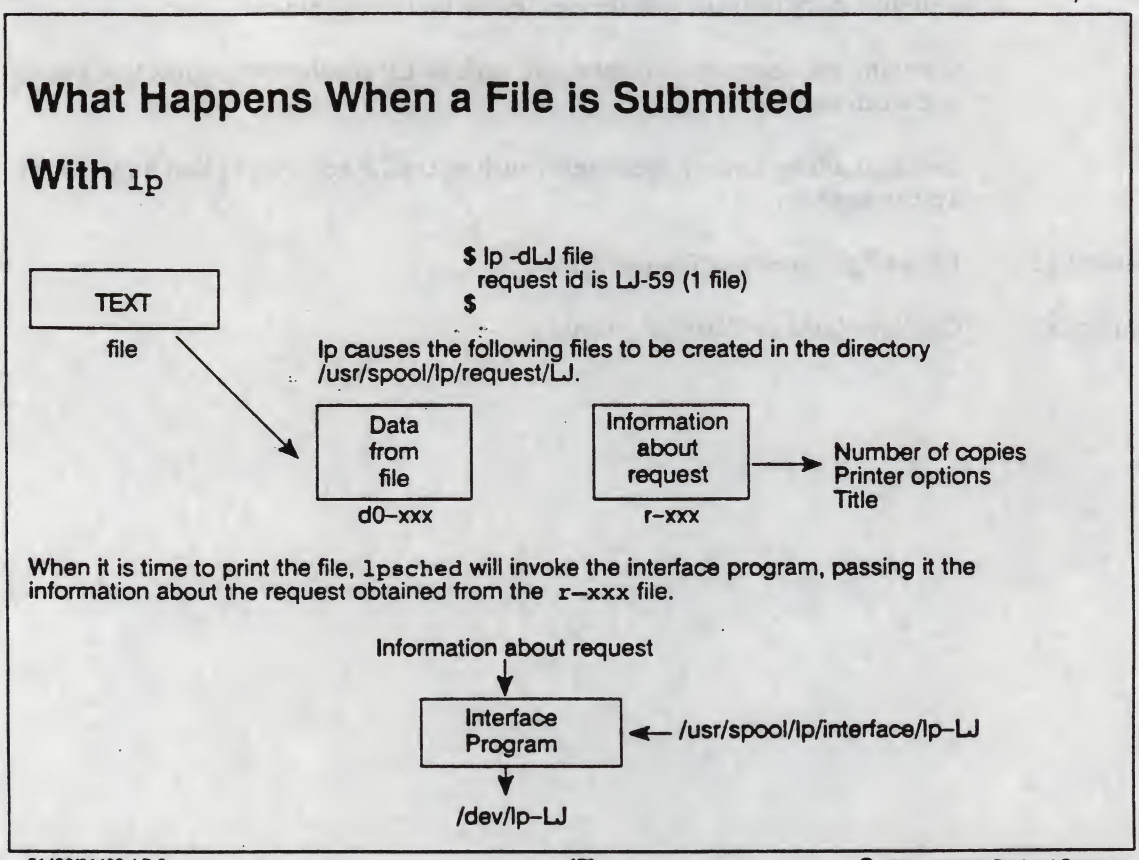


51436/51482 LP-3. 170 

### Student Notes

When the `lp` command is invoked, a print destination is determined and a print request is submitted to the spooling system. A unique request id number is assigned to each print request.

The print destination (in our example, LJ) is a logical name known to the spooling system and related to a spool directory. Usually, the destination is the name of a particular printer. However, several printers can be combined in to a class of printers, so a class may also be a destination. If the destination of a print request is the name of a printer, the request will go to that printer's queue only; however, if the destination is a class of printers, the request will be spooled on to that class's queue and will be printed by the first printer that is available within that class. This helps avoid backups in a particular queue.

All print requests are queued in their destination spool directory strictly in priority order, that is the job with the highest priority will print first. If two jobs have the same priority they will be printed in FIFO (first-in-first-out) order.

# Module LP — Managing the LP Spooler

Priorities are set both for printers and print jobs. The printer priority is set when you set up the printer, either with SAM or the `lpadmin` command, although it can be changed on the fly. Print jobs are given priority when they are submitted with the lp command. These jobs can also have their priority modified if desired.

What happens when a job is submitted is that its priority is checked against that of the printer that it was sent to. If its priority is higher than or equal to that of the printer it will be printed, but not before other jobs of higher priority. If its priority is below that of the printer then the job will not be printed until either the job's priority is raised or the printer's priority is lowered, or the job is moved to another printer with lower priority.

Before a print request is actually printed, an interface script is invoked by the spooling system. An **interface script** is the link between the scheduler and the device (that is, its device file). It takes arguments from the scheduler and sends header and configuration information to the device. Each device has its own interface script which can be modified by the administrator.

bv. lp -dlp -o landscape.

wordt aan script in /lp/interface meegegeven.

-o raw.

bv. voor images; wordt niet 'vertaald'.

-onb

no banner.

## LP-4. SLIDE: LP Spooler Administrative Commands

184

### LP Spooler Administrative Commands

- Scheduling

```
# lpshut
# lpsched
```

stopper
start

- Queueing

```
# reject -r"reason" printer_name
# accept printer_name
```

- Printing

```
# disable -r"reason" printer_name
# enable printer_name
```

- Reporting spooler status

```
# lpstat -t
```



## Student Notes

The commands on the slide are explained in more detail below.

```
# lpshut
```

**lpshut** stops the scheduler, which in turn stops all printers. Internally, the scheduler is a daemon process that sends files from destination queue directories to the interface scripts. Whenever the structure of the spooling system needs to be changed (for example, by commands such as **lpadmin** or **lpmove**, the scheduler must be stopped. All requests printing when **lpshut** is executed will be stopped and left in the queue. When the scheduler is started again, it will begin printing those files that were stopped as if they were a new request.

Note, even though the scheduler is in a halted state, **lp** can still queue jobs to the various destinations.

```
# lpsched
```

`lpsched` starts the scheduler. It is normally invoked from the `/etc/rc` file during the boot-up procedure. `lpsched` is the opposite of `lpshut` and must be executed after modifying the spooling system structure in order to start proper operation again.

```
# reject -r"reason" printer_name
```

`reject` temporarily rejects requests to printer_name. The `lp` command is prevented from queuing further print requests to this destination. If a user executes `lp` to the destination printer_name, the user will get an error message similar to this:

```
lp: can't accept requests for destination "printer_name"
```

The `accept` command is the opposite of `reject`. It changes the status of a destination to allow the `lp` command to queue requests to this destination. This command is executed when the destination is first created. It need not be executed again unless the LP administrator decides that this destination should not queue up requests, for whatever reasons, and executes a `reject`.

```
# disable -r"reason" printer_name
```

This command disables the scheduler from sending request to the printer. Users, however, can still submit requests with the `lp` command.

`enable is the opposite of disable.` `enable` allows the scheduler to send request to printer_name.

```
lpstat -t
```

`lpstat` allows you to check the status of the spooling system. `lpstat -t` prints the status of the scheduler and each printer configured in the spooler. For each printer, it tells you if that printer is enabled or disabled, is accepting or rejecting requests, and the request-id of the files that are queued for printing.



27c

## LP-5. Other Things the Administrator Can Do

p85

### Other Things the Administrator Can Do

[illegible] lpshut

- Insert the printer "laserjet" into the class of printers "fastprint"

  `# lpadmin -plaserjet -cfastprint`

- Remove a printer "qj1" from the class of printers "regprint"

  `# lpadmin -pqj1 -rregprint`

- Move all requests from destination "lp1" to destination "laserjet"

  `# lpmove lp1 laserjet`

- Set the system default printer

  `# lpadmin -dlaserjet`

51436/51482 LP-5. 172 

### Student Notes

The commands on the slide are explained in more detail below.

```
# lpadmin -plaserjet -cfastprint
```

This command inserts the printer "laserjet" into the class of printers "fastprint". Thereafter print requests can be addressed to the destination "fastprint". Note that the scheduler must have been stopped before invoking this `lpadmin` command.

```
# lpadmin -pqj1 -rregprint
```

This command removes the printer "qj1" from the class of printers "regprint" (for "regular printers"). If "qj1" was the last member of the destination "regprint", the class itself would be deleted.

```
# lpmove lp1 laserjet
```

This command moves all requests from destination "lp1" to destination "laserjet". Furthermore, a `reject` command is internally executed for the destination "lp1". If one printer is out of order, for example, due to hardware failure, and there are still print requests waiting to be printed on that specific printer, you may want to move the jobs to another printer that is operating. Note that the scheduler must be stopped before executing `lpmove`.

```
# lpadmin -dlaserjet
```

This command makes "laserjet" the new system default destination. "laserjet" must already be an existing destination. The `-d` option may be used when `lpsched` is running.

```
# lpadmin -plaserjet -g4
```

This command changes the priority value on the printer "laserjet" to 4. This means that only jobs scheduled with a priority higher then 4 will print on this printer. The range of priority for printers and print jobs is 0 to 7, 0 being the lowest and the default and 7 being the highest. Users would use the -p option with the `lp` command to submit jobs with a certain priority:

```
$ lp -p3 file1
```

This example submits a job to the default printer with priority 3.

---

**Note** Always shutdown the scheduler before using the `lpadmin` command.

---



271

## LP-6. How to Add a Printer to the LP Spooler

186

### How to Add a Printer to the LP Spooler

Before beginning, make sure the printer is connected and you have made a device file.

1. Shut down the LP scheduler.

   ```
   # lpshut
   ```

2. Check /usr/spool/lp/model and select the appropriate model script.

   priviliy

3. Add the printer to the LP spooler. modelscript

   ```
   # lpadmin -plaserjet -v/dev/lp1 -mPCL3 -cclass1 -g4
   ```

4. Allow printer to accept print requests.

   ```
   # accept laserjet
   ```

5. Enable the printer to print the requests.

   ```
   # enable laserjet
   ```

6. Restart the LP scheduler.

   ```
   # lpsched
   ```

51436/51482 LP-6. 173 

### Student Notes

Before adding your printer to the LP spooler system, make sure it is connected and the appropriate device file exists. Creating the appropriate device file is different if you are on a Series 300 versus a Series 800. To create the device file on a Series 800, use the **mksf** command. For example, to create a device file for a printer on mux0, port 1:

```
# mksf -d mux0 -l o -p 1 -h /dev/lp1
```

To create the device file on a Series 300, use the **mknod** command. For example, to create a device file for a printer at select code 9 on port 1:

```
# mknod /dev/lp1 c 7 ox090104
```

We will cover device files in more detail later.

# Module LP — Managing the LP Spooler

Once you have connected your printer to your system and have created the appropriate device files, you can add the printer to the LP spooler system. The steps to do this are shown on the slide. (The commands referenced on the slide are located in the `/usr/lib` directory.)

Once you choose the model script from the `/usr/spool/lp/model` directory, make sure the script has:

- A permission mode of 644 (-rw-r—r—)
- Is owned by `lp`
- Is in the group `bin`

The `lpadmin` command configures the LP spooler to describe printers, classes, and devices. It is used to add and remove destinations, change membership in classes, change devices for printers, change printer interface programs, and to change the system default destination.

The `lpadmin` example shown on the slide adds the printer to the LP spooler system. The command is shown below and each of the options is described in more detail.

```
# lpadmin -plaserjet -v/dev/lp1 -mPCL3 -cclass1 -g4
```

The options mean the following:

| | |
|---|---|
| -p*printer* | Names a printer to which all of the options refer. |
| -v*device* | Specifies the full pathname of the printer's device file. Note there is nothing to stop an administrator from associating the same *device* with more than one printer. |
| -m*model* | Selects a model interface program for the printer. *model* is one of the model interface names in the `/usr/spool/lp/model` directory. |
| -c*class* | Inserts the printer into the specified *class*. This is optional; printers do not have to belong to a class. |
| -g*priority* | Sets the default priority for the printer. This is optional; the default is 0. |

Note that when you use the command `lpadmin`, exactly one of the following options must be supplied:

| | |
|---|---|
| -p*printer* | Names a printer to which all of the options refer. |
| -x*dest* | Removes destination *dest* from the LP system. |
| -d*dest* | Makes *dest*, which already exists as the default destination. |

There are many more options to the `lpadmin` command. See *LPADMIN(1M)* for more information.

Once you restart the scheduler with the `lpsched` command, you can verify that the scheduler is functioning with the `lpstat -t` command.

```
# lpstat -t
scheduler is running
system default destination: lp
device for lp: /dev/lp
device for laserjet: /dev/lp1
lp accepting requests since Aug 17 10:29
laserjet accepting requests since Sep 10 11:15
printer lp is idle.  enabled since Aug 17 10:29
     fence priority: 0
```



272

```
printer laserjet is idle.  enabled since Sep 10 11:15
     fence priority: 0
```

**Note** If you are adding a serial printer to your system, make sure there is not a getty running on the serial port.

## LP-7. The SCHEDLOCK File

187

### The SCHEDLOCK File

- lpsched creates the file SCHEDLOCK when it starts the scheduler
- lpshut removes the files SCHEDLOCK when it shuts down the scheduler
- lpsched will only allow one scheduler to run at a time; when invoked, it checks for the existence of SCHEDLOCK to see if a scheduler is currently running; if SCHEDLOCK exists, lpsched will not start up the scheduler
- If the scheduler won't start:

```
# rm -f /usr/spool/lp/SCHEDLOCK
```

- If the scheduler won't stop using lpshut:

```
# kill -15 lpsched_pid
```

51436/51482 LP-7. 174 

## Student Notes

If, for some reason, the scheduler won't start up when you execute the **lpsched** command, check to see if the SCHEDLOCK files exists in the **/usr/spool/lp** directory. If it exists, remove it with this command:

```
rm -f /usr/spool/lp/SCHEDLOCK
```

Then try again to start the scheduler with **lpsched**.

SCHEDLOCK prevents the execution of more than one scheduler. So, if the SCHEDLOCK file exists when you try to start the scheduler running, the **lpsched** command won't start up the scheduler. This is why you should remove the file and then try again to start the scheduler.

The **lpshut** command automatically removes the SCHEDLOCK file when it terminates the scheduler. If the SCHEDLOCK file does not exist when you execute **lpshut**, then **lpshut** might not work. If **lpshut** won't shut down the scheduler, you can use the following command:



273

```
# kill -15 lpsched_pid
```

When you start the scheduler again with `lpsched`, it will create the SCHEDLOCK file.

## LP-8. SLIDE: SAM Versus Manual Commands

188

### SAM Versus Manual Commands

```
SAM                          Printers and Plotters

        The choices below pertain to plotters as well as printers. Highlight an
        item and then press "Return" or "Select Item."

                Printers:
                    Add a Local Printer ...
                    Add a Remote Printer ...
                    Enable a Printer ...
                    Disable a Printer ...
                    Remove a Printer ...

                LP Spooler Administration:
                    Set the System Default Printer ...
                    Start Up (Shut Down) the Spool System
                    View Printer Status Information
                    View Print Requests

 Help    Main Menu    Shell    Select Item                           Previous Menu
```

51436/51482 LP-8. 175 

### Student Notes

SAM can be used as well as manual commands to administer the LP spooler system. The tasks that SAM can perform are shown on the slide. This menu will come up after choosing:

SAM Main Menu | Peripheral Devices -> | Printers and Plotters ->

In general, SAM is easier to use than the manual commands. You are led through each task with menus and templates. You are prompted for all necessary entries, and the entries are checked for plausibility. If an entry isn't correct, you get an error message and are prompted for a different response.

Though you may find SAM easier to use, it is less flexible than the manual commands. You cannot control many of the options available with commands. Also, there are some commands that are not available in SAM. For example, SAM cannot:

- submit a request
- move a request between printers

- cancel a request
- accept and reject requests

## LP-9. SLIDE: Adding a Printer to the LP Spooler System With SAM

189

### Adding a Printer With SAM (300/400):

```
SAM                         Add a Local Printer

        Fill in or modify the desired fields and then press "Perform Task".

        Printer name . . . . . . . . . . . . . . . . . . _______

        Printer model/interface  . . . . . . . . . . . __

        Printer device file name . . . . . . . . . . . __

        Printer priority (0 - lowest, 7 - highest) . . 0

        Make this the system default printer? (y or n) n

        Printer class . . . . . . . . . . . . . . . . . _________ (optional)

        Printer connected to a terminal? (y or n) . . . _ (optional)

[Help] [Main Menu] [Shell] [Perform Task] [ ] [ ] [ ] [Exit Task]
```

51436/51482 LP-9. 176 

### Student Notes

To add a printer with SAM, choose the following menus:

```
System Administration Manager
          |
          V
Peripheral Devices ->
          |
          V
Printers and Plotters ->
          |
          V
Add a Local Printer ...
```

After you select "Add a Local Printer," you will get the screen shown on the slide. Fill in the information at the blanks.



275

- Name your printer. This is the name that will be recognized by the LP Spooler system.
- Give a model script. If you are not sure what script to use, press the (Help) key. SAM will give you a list of the interface scripts in the `/usr/spool/lp/model` directory.
- Enter the device file name. If the device file does not exist, SAM will ask you if you want it to be created. (If you want SAM to create the device file, it will ask you for more information after you press the (Perform Task) key below.)
- Everything else is optional. Fill in the information if you wish.
- Press the (Perform Task) key and SAM will add the printer to LP Spooler system.

In the example, we are adding an HP QuietJet PLUS printer (with an RS-232C interface) to the spooling system. We named the printer "`qj1`" (as a acronym for QuietJet #1). To select the model/interface, we chose the one which matches the product number of the QuietJet printer, `quietjet`. The printer is attached to the sixth port on the first MUX board, so we chose the device file name `/dev/lp5`.

If a problem occurs, you can escape to a shell by pressing (shell) and try to fix the problem.

## LP-10. WORKSESSION: Review Questions

1. What functions does the spooling system provide, and why are they required?

2. Which of these functions are available to the administrator and which to normal users?

3. What is the difference between a device, a printer and a destination?

4. Is it possible, to install two printers for one device?

5. How would you cancel your own print request?

6. How would you cancel a print request owned by someone else?

7. How would the owner of a canceled request know that you had canceled it?

8. What is an interface program?

9. If you have stopped the scheduler (by `lpshut`), does print-out continues?

10. If you have stopped the scheduler (by `lpshut`), can you still use the `lp` command to add print requests to the queues?

11. How can you tell other users that a printer is "broken"?

12. How can you redirect print requests from the "printer1" queue to the "printer2" queue?

# Module LP — Managing the LP Spooler

## LP-11. LAB: Lab Exercises (Series 300 Only)

### Directions

Work with your lab partner to complete the following exercises.

1. Physically connect your line printer to your system (if it is not already connected.)

2. Set up the LP spooler so that the command `lp /etc/passwd` prints the password file to your printer.

3. Check the status of your line printer spooler.

4. Disable the printer. Now send some output to the printer. Determine if the files are queued. Now cancel one of the requests from the queue. Enable the printer. What happens?



277

## LP-12. LAB: Lab Exercises (Series 800 only)

### Directions

In this lab, you are going to work in pairs of terminals. Your instructor will assign you and your partner a terminal. You should determine the device file for the terminal that you will use as a pseudo-printer. You can find this device file by doing a `who am i` on the terminal to be used as a printer. Step 9 will be done by the instructor when everyone has finished steps 1 through 8. For steps 6 and 11, you need the token. Don't do this either of these steps unless you have the tolken.

1. Determine the status of the LP scheduler. It should not be running.

2. Add a printer with the `lpadmin` command by specifying the name, device file, and model (use any model name).

3. Since there can be only one default destination, assign the default destination to the printer specified by the instructor. (This only needs to be done by one group, but everyone should check it.)

4. Allow the scheduler to accept requests for your new printer.

5. Enable the scheduler to print to your new "printer."

# Module LP — Managing the LP Spooler

6. Edit the file `/etc/inittab` and comment out the line which runs a `getty` on the device which you are going to use as a printer.

7. Kill the `getty` currently running on your "printer."

8. Run a `sleep` command (nohuped) in the background on the terminal port. This will allow the `stty` command in the model to assign a reasonable protocol for the serial port.

The syntax is: `nohup sleep 600000 < /dev/`*tty_port* `&`

9. The instructor should turn the scheduler back on.

10. Try printing to your terminal with the `lp` command. You will need to check out the `-d` option with the `man` command.

11. When you are done, kill your `sleep` command. Your instructor will edit the `/etc/inittab` file and uncomment the `getty` lines. Your instructor will then shutoff the scheduler. You should remove your printer from the scheduler. When everyone is ready, the instructor will start the scheduler and invoke `init 2` to restart all the gettys.



278

# Module LP — Managing the LP Spooler